TANGLE STRUCTURE TREES

HANNO VON BERGEN* AND REINHARD DIESTEL*

Abstract. We introduce a comprehensive data structure, tangle structure trees, which simultaneously displays all the $\mathcal{F}$-tangles of a separation system – e.g., the separations of a graph or the bipartitions of a large data set – for very general obstruction sets $\mathcal{F}$. It simultaneously also displays certificates $\sigma \in \mathcal{F}$ for any non-existence of such tangles, or for the non-extendability of low-order tangles to higher-order ones. Our main result characterizes the types of tangle, in any separation system, for which such structure trees exist.

Our result can be applied to produce the structures of the classical tree-of-tangles and tangle-tree duality theorems, both for graph tangles and for their known generalizations to more general separation systems. It extends those theorems to obstruction sets $\mathcal{F}$ that need not define profiles (as they must in all known trees of tangles) or consist of stars of separations (as they must in traditional tangle-tree duality theorems).

Our existence proof for these structure trees is constructive. The construction has been implemented in open-source software available for tangle detection and further analysis [1].

1. Introduction

The notion of 'tangles' was originally introduced by Robertson and Seymour [20] as an abstract concept of high local connectivity in graphs, one that unifies several more concrete such notions, such as highly connected subgraphs or minors. What all these notions have in common is that, given any low-order separation of the graph, any highly cohesive substructure must lie mostly on one of its two sides: since the separation has low order, it cannot split it into two roughly equal halves. A tangle remembers only how all the low-order separations are oriented in this way, each towards that highly cohesive substructure: the collection of all these oriented separations is called a *tangle*. See [11] for a precise definition and basic facts about graph tangles. Algorithmic aspects of graph tangles, specifically of our theoretical results in [7], were first studied by Grohe and Schweitzer [17].

Tangles have since been generalized to more general settings than graphs. In the setting of set partitions, they offer a precise theoretical basis for 'fuzzy' real-world problems such as clustering in large data sets [10, 19]. All these, including graphs, are special cases of so-called *abstract separation systems* [9]: an axiomatic setting that assumes only the most basic properties of graph separations. This is the setting in which tangles are most easily and comprehensively treated, and we shall use this framework also in this paper.

There are two fundamental theorems about tangles, which both have their origins in [20] and are also treated in [11]. The *tree-of-tangles* theorem exposes a tree-like structure in the graph or data set whose tangles we consider, which 'distinguishes' the tangles in that they are shown to live in different areas of this tree-like structure.

Extended Abstract for STOC.

*) Fachbereich Mathematik, Universität Hamburg.

The other is the *tangle-tree duality* theorem, which certifies the non-existence of a tangle by exposing another tree-like structure in the graph or data set, one whose existence clearly precludes the existence of a tangle.[1] This is the more fundamental of the two theorems, in the sense that the other can be reduced to its abstract version (see below) but not vice versa [15].

Both the above theorems have been generalized to $\mathcal{F}$-tangles of abstract separation systems, ways of simultaneously orienting the 'separations' of a given structure in such a way that no three of those oriented separations form an element of $\mathcal{F}$.[2] This collection $\mathcal{F}$ has to satisfy some constraints that depend on which of the two theorems we are considering.

For the tree-of-tangles theorem, which displays all the $\mathcal{F}$-tangles of a separation system simultaneously, $\mathcal{F}$ has to include triples of oriented separations that are reminiscent of ultrafilters. If our separations are bipartitions of a fixed set V, for example, then $\mathcal{F}$ must contain all triples of the form $\{A, B, \overline{A \cap B}\}$: if A and B are subsets of V deemed 'large' (in that the $\mathcal{F}$-tangle orients those partitions towards them), then the complement of $A \cap B$ cannot also be 'large'. Such $\mathcal{F}$-tangles are called *profiles*, and all existing tree-of-tangle theorems for $\mathcal{F}$-tangles assume that these are profiles [2, 3, 5, 6, 8, 12, 13, 16].

For the tangle-tree duality theorem, $\mathcal{F}$ has to satisfy another constraint: its elements have to be *stars*, nested sets of oriented separations pointing towards each other [9]. The tree-like structures by which tangle-tree duality theorems certify the non-existence of $\mathcal{F}$-tangles require that $\mathcal{F}$ consist of stars: if it does not, these structures cannot be tree-like.

This state of the art leaves $\mathcal{F}$-tangles that are not profiles without any known way of displaying all those tangles simultaneously. And if $\mathcal{F}$ does not consist of stars, it leaves separation systems that have no $\mathcal{F}$-tangles without any known way of organizing the elements of $\mathcal{F}$ into a data structure that displays them as easily checkable certificates for the non-existence of such tangles.

The structure trees whose existence we prove in this paper do both these things: they display all the $\mathcal{F}$-tangles of a separation system even when they are not profiles, yielding more general tree-of-tangles theorems [4, Theorems 6.2 & 6.10], and for those orientations of the separation system that are not $\mathcal{F}$-tangles they display certificates in $\mathcal{F}$ that show why they are not. If $\mathcal{F}$ happens to consist of stars, our results yield the traditional certificates from a new premise.

Tangle structure trees display all this information in a single, comprehensive, data structure, which is maximally efficient in the following structural sense. A single tangle is most efficiently displayed by listing just those of its elements that are minimal in the poset of oriented separations that comes with a separation system. Indeed, all these are needed to determine the tangle, but any other elements can be deduced from them. Our structure trees display all the $\mathcal{F}$-tangles of a separation system S simultaneously by listing only the oriented separations that are minimal in one of those tangles, including the tangles of subsystems consisting of separations of lower order. And for orientations of S that are not $\mathcal{F}$-tangles it displays certificates from $\mathcal{F}$ inside those same minimal subsets of oriented separations, in a way readily accessible in the structure tree.

[1] In a graph, this would be a tree-decomposition into parts too small for a tangle to live in.

[2] In the case of graph tangles, $\mathcal{F}$ consists of the sets of up to three oriented separations of the graph such that the union of the 'back' sides of these three separations covers the entire graph [11].

2. Tangle basics

In this section we give precise definitions and notation for abstract tangles, largely following [9] and indicating any deviations.[3]

Tangles of graphs are ways of orienting their separations, each towards one of its two sides. Abstract tangles are designed to work in scenarios where there need not be anything to 'separate'. In order to retain our intuition from graphs, however, we continue to refer to the things of which our abstract tangles pick one of two variants (which they will indeed do) as 'separations'. These are defined by noting some key properties of graph separations and making them into axioms, as follows.

A *separation system* $(\vec{S}, \leqslant, ^*)$ is a set $\vec{S}$, whose elements we call *oriented separations*, that comes with a partial ordering $\leqslant$ on $\vec{S}$ and an order-reversing involution $^*\colon \vec{S} \to \vec{S}$. Thus, for any two elements[4] $\vec{r}, \vec{s}$ of $\vec{S}$ with $\vec{r} \leqslant \vec{s}$ we have $\vec{r}^* \geqslant \vec{s}^*$. We write $\vec{s}^* =: \overleftarrow{s}$, and call $\overleftarrow{s}$ the *inverse* of $\vec{s}$. While we allow formally that $\vec{s} = \overleftarrow{s}$, in which case we call $\vec{s}$ and s *degenerate*, this does not happen often in practice.[5]

If a separation system $\vec{U}$ happens to be a lattice, that is, if there is a supremum $\vec{r} \vee \vec{s}$ and an infimum $\vec{r} \wedge \vec{s}$ in $\vec{U}$ for every two elements $\vec{r}, \vec{s} \in \vec{U}$, we call $\vec{U}$ a *universe* of separations. It is *distributive* if it is distributive as a lattice. A separation system $\vec{S} \subseteq \vec{U}$ is *submodular* if for every two elements of $\vec{S}$ either their infimum or their supremum in $\vec{U}$ also lies in $\vec{S}$.

Very rarely we may have separations $\vec{s} \leq \overleftarrow{s}$; then $\vec{s}$ is *small* and $\overleftarrow{s}$ is *large*.[6] We say that $\vec{s}$ is *trivial* (and $\overleftarrow{s}$ is *co-trivial*) in $\vec{S}$ if there exists a pair of inverse separations $\vec{r}, \overleftarrow{r} < \vec{s}$ in $\vec{S}$. Trivial separations are clearly large, so co-trivial ones are small, but the converse need not hold. See [9] for more on these technicalities if desired.

The set of *unoriented separations* in $(\vec{S}, \leq, ^*)$ is

$$S := \{\{\vec{s}, \overleftarrow{s}\} : \vec{s} \in \vec{S}\}.$$

We call the elements $\vec{s}, \overleftarrow{s}$ of s its *orientations*. An *orientation of* S is a set $\tau \subseteq \vec{S}$ that contains exactly one orientation of every $s \in S$. For $s \in S$ we then denote by $\tau(s)$ the unique orientation of s contained in τ. An orientation of a subset of S is a *partial orientation* of S.

If $\vec{r} \geqslant \vec{s}$ we say that $\vec{r}$ *points towards* s (and that $\overleftarrow{r}$ *points away from* s). We say that $\vec{r}$ *points towards* an oriented separation $\vec{s}$ whenever it points towards s, i.e., if $\vec{r} \geqslant \vec{s}$ or $\vec{r} \geqslant \overleftarrow{s}$, and similarly for 'points away from'. A *star* is a set σ of non-degenerate oriented separations that point towards each other. As is easy to check, this happens if and only if $\vec{r} \geqslant \overleftarrow{s}$ (and hence $\vec{s} \geqslant \overleftarrow{r}$) for all distinct $\vec{r}, \vec{s} \in \sigma$.

Two separations $s, r \in S$ are *nested* if they have orientations that are comparable under $\leqslant$. Oriented separations are *nested* if their underlying unoriented separations are nested. A subset of $\vec{S}$ is *nested* if its elements are pairwise nested.

A subset of $\vec{S}$ is *consistent* if no pair of its elements $\vec{r}, \vec{s}$ with $r \neq s$ point away from each other. Stars are examples of consistent nested sets of oriented separations.

---

[3] The most important difference is that, for historical reasons, the partial ordering on $\vec{S}$ used in [9] is the inverse of ours. So terms like 'large' and 'small', infima and suprema etc, are reversed.

[4] We often denote the elements of $\vec{S}$ by letters with an arrow, in either direction, precisely in order to have a simple way to refer to their dual elements: by reversing the arrow. But the arrow directions have no meaning: an arbitrary element of $\vec{S}$ could be denoted equally as $\vec{s}$ or as $\overleftarrow{s}$.

[5] The only degenerate separation of a graph $G = (V, E)$, for example, is $\{V, V\}$.

[6] The small separations of a graph G are those of the form (V, A) with $A \subseteq V = V(G)$.

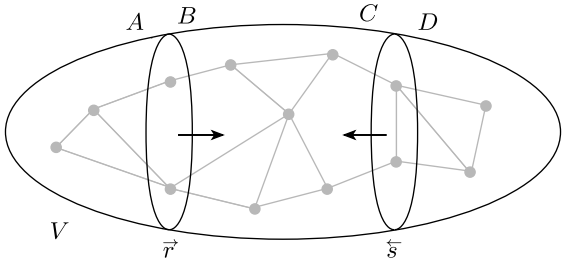


FIGURE 1. Nested separations $r = \{A, B\}$ and $s = \{C, D\}$ of a graph. Their orientations $\vec{r} = (A, B)$ and $\overleftarrow{s} = (D, C)$ point towards each other, since $\vec{r} \geq \vec{s}$ (as $B \supseteq D$) and $\overleftarrow{s} \geq \overleftarrow{r}$ (as $C \supseteq A$).

We shall often be interested in consistent orientations of S. For each of its elements $\vec{r}$, a consistent orientation of S will also contain every $\vec{s} > \vec{r}$ other than, possibly, $\vec{s} = \overleftarrow{r}$. Consistent partial orientations of S are easily seen to extend to consistent orientations of S, unless they contain a separation that is co-trivial in $\vec{S}$; see [9].

The *closure* in $\vec{S}$ of a consistent set $\sigma \subseteq \vec{S}$ is the set

$$\lfloor \sigma \rfloor = \sigma \cup \{ \vec{s} \in \vec{S} : \exists \vec{r} \in \sigma \text{ such that } r \neq s \text{ and } \vec{s} > \vec{r} \}.$$

Pathological exceptions aside, $\lfloor \sigma \rfloor$ is again consistent, and $\lfloor \lfloor \sigma \rfloor \rfloor = \lfloor \sigma \rfloor$.

An *order function* on S is any map $S \to \mathbb{R}$. Unless otherwise mentioned, we denote such order functions as $s \mapsto |s|$. We extend them to $\vec{S}$ by letting $|\vec{s}| := |\overleftarrow{s}| := |s|$. A separation system $(\vec{S}, \leq, ^*)$ given with an order function on S is an *ordered separation system*. For every $k \in \mathbb{R}$ we let $\vec{S}_k := \{\vec{s} \in \vec{S} : |s| < k\}$. We say that $\tau \subseteq \vec{S}$ *avoids* a set $\mathcal{F} \subseteq 2^{\vec{S}}$ of *forbidden subsets* if τ has no subset in $\mathcal{F}$.

Definition 2.1. An *$\mathcal{F}$-tangle* of S is an $\mathcal{F}$-avoiding consistent orientation of S. The $\mathcal{F}$-tangles of the subsets S_k of S are the $\mathcal{F}$-tangles *in* $\vec{S}$.

3. Tangle structure trees

We adopt the graph-theoretic terminology of [11]. A *tree* is a connected acyclic graph. Given nodes t, t' of a tree T we write tTt' for the unique t–t' path in T. A *rooted tree* is a tree with a distinguished node called its *root*. Given a tree T with root r, we define a partial order $\leqslant_r$ on $V(T)$ by declaring $x \leqslant_r y$ if x lies on the path in T from r to y. Maximal elements, including the root if $|T| = 1$, are called *leaves*. Any direct successors in $<_r$ of a node of T are its *children*. We write E_v for the set of edges of T from a node v to its children.

Let $(\vec{S}, \leq, ^*)$ be a separation system, and let $\mathcal{F}$ be any set.

Definition 3.1. A *separation tree* (T, r, β) on $\vec{S}$ consists of a rooted tree (T, r) together with an edge labelling $\beta \colon E(T) \to \vec{S}$ such that for every non-leaf $v \in V(T)$ there exists a separation $s_v \in S$ such that β restricts to a bijection $E_v \to \{\vec{s_v}, \overleftarrow{s_v}\}$ and $s_u \neq s_v$ whenever $u <_r v$.

Thus, every non-leaf node v of a separation tree has either one or two children. If there is no need to refer to r or β explicitly, we usually abbreviate (T, r, β) to T. For every node v, we write β_v for the set $\beta\big(E(rTv)\big)$ of edge labels on the path rTv. A separation tree T is *consistent* if $\beta_v \subseteq \vec{S}$ is consistent for every node $v \in T$.

Consistent separation trees offer an efficient data structure for storing consistent orientations of S, in particular, for tangles. Indeed, any such orientation τ labels

the tree's edges between the root and some leaf ℓ, but this β_ℓ is only a small subset of τ: elements of τ in $\lfloor \beta_v \rfloor$ for some $v < \ell$ will not appear in $\beta_\ell \setminus \beta_v$. But they can be reconstructed from β_v by the consistency of τ.

A leaf ℓ of a consistent separation tree on $\vec{S}$ is an ($\mathcal{F}$-) *tangle leaf* if the closure $\lfloor \beta_\ell \rfloor$ of β_ℓ is an $\mathcal{F}$-tangle of S. A non-leaf node v is a *tangle node* if there is a tangle leaf $\ell \geq_r w$ for every successor w of v in $>_r$. A leaf ℓ is *forbidden* (by $\mathcal{F}$) if β_ℓ contains an element of $\mathcal{F}$ as a subset. Note that tangle leaves are never forbidden.

Definition 3.2. An *$\mathcal{F}$-tangle structure tree of $\vec{S}$* is a consistent separation tree on $\vec{S}$ in which every leaf is either a tangle leaf or forbidden, and for every non-leaf node v the set β_v has no subset in $\mathcal{F}$.

Theorem 3.3. *Let $\vec{S}$ be a separation system, let $\mathcal{F}$ be any set, and let T be any $\mathcal{F}$-tangle structure tree of $\vec{S}$.*

(i) *For every $\mathcal{F}$-tangle τ of S there is a unique leaf ℓ of T such that $\lfloor \beta_\ell \rfloor = \tau$.*
(ii) *If all the leaves of T are forbidden, then S has no $\mathcal{F}$-tangle.*

4. Existence of tangle structure trees

Let $\vec{S}$ be an ordered separation system, and let $\mathcal{F}$ be any set. We shall need two conditions on $\mathcal{F}$ to ensure that $\vec{S}$ has a tangle structure tree.

The first is that $\{\overleftarrow{s}\} \in \mathcal{F}$ for every $\vec{s} \in \vec{S}$ that is trivial in $\vec{S}$. This holds for all sets $\mathcal{F}$ of interest, and if it does we call $\mathcal{F}$ *standard for* $\vec{S}$.

The second condition has more substance. If T is any separation tree on $\vec{S}$, then every orientation τ of S contains β_ℓ for a unique leaf ℓ. If τ is consistent and $\lfloor \beta_\ell \rfloor$ orients all of S, then $\lfloor \beta_\ell \rfloor = \tau$. If τ is not a tangle, then $\tau = \lfloor \beta_\ell \rfloor$ has a subset σ in $\mathcal{F}$ that witnesses this. If T is a tangle structure tree, we can find such a set $\sigma \in \mathcal{F}$ not only in $\lfloor \beta_\ell \rfloor$ but even in β_ℓ: among the edge labels of T.

In order for this to be possible, we therefore need to make some 'richness' assumption about $\mathcal{F}$: an assumption which, in our example, ensures that $\mathcal{F}$ has enough elements to contain a subset also of β_ℓ as soon as it contains a subset of $\lfloor \beta_\ell \rfloor$. We shall identify an essentially weakest-possible such richness condition below.

A separation tree T on an ordered separation system $\vec{S}$ is *ordered* if $|s_v| \leq |s_w|$ whenever v and w are non-leaves of T with $v \leq w$. It is *thoroughly ordered* (in $\vec{S}$) if, for every non-leaf node v, the separation s_v is not oriented by $\lfloor \beta_v \rfloor$ and has minimum order among the separations in S not oriented by $\lfloor \beta_v \rfloor$.

We say that $\vec{s} \in \vec{S}$ is *weakly eclipsed* by $\vec{r} \in \vec{S}$ if $\vec{r} < \vec{s}$ and $|r| \leq |s|$, and *eclipsed* by $\vec{r}$ if $\vec{r} < \vec{s}$ and $|r| < |s|$. Given any set $\tau \subseteq \vec{S}$, a subset $\sigma \subseteq \tau$ is *efficient* (in τ) if no element of σ is eclipsed by any other element of τ. It is *strongly efficient* if no element of σ is weakly eclipsed by any other element of τ. Note that if the order function on S is injective and τ is a partial orientation of S, then every efficient subset of τ is strongly efficient in τ.

Definition 4.1. A set $\mathcal{F}$ is *rich for* $\vec{S}$ if every consistent orientation of S that has a subset in $\mathcal{F}$ also has a strongly efficient[7] subset in $\mathcal{F}$.

Although it may look a bit technical at first glance, this 'richness' requirement on $\mathcal{F}$ is quite natural, given the role of these $\mathcal{F}$ in tangle theory. For example, given a consistent orientation τ of S that has a subset $\sigma \in \mathcal{F}$, we can make it strongly

[7] in this orientation of S

efficient by replacing every element $\vec{s}$ of σ by some $\vec{s'} \leq \vec{s}$ that is minimal in τ. If the sets σ' obtained in this way are always again in $\mathcal{F}$, then $\mathcal{F}$ will be rich. And for many natural $\mathcal{F}$ they will be. For the idea behind the sets in $\mathcal{F}$ is that they identify areas in our graph or other structure that are 'too small to be home to a tangle'. Hence it is not unreasonable for σ' to be in $\mathcal{F}$ if σ was, given the role played by $\mathcal{F}$.

The assumption that $\mathcal{F}$ is rich is central to our results. It is needed for $\mathcal{F}$-tangle structure trees to exist, and no weaker condition on $\mathcal{F}$ will ensure the same:

Theorem 4.2. *Let $\vec{S}$ be an ordered separation system, and let $\mathcal{F}$ be standard for $\vec{S}$. If $\mathcal{F}$ is rich, there exists a thoroughly ordered $\mathcal{F}$-tangle structure tree of $\vec{S}$. If the order function on S is injective, this $\mathcal{F}$-tangle structure tree is unique if it exists, and its existence implies that $\mathcal{F}$ is rich.*

5. Irreducible and efficient tangle structure trees

The thoroughly ordered tangle structure trees whose existence we established in Section 4 can be quite large. But we can shrink them in a way that maintains their dual use for displaying both tangles and certificates for their non-existence.

Let (T, r, β) be a consistent separation tree on a separation system $\vec{S}$. Let v be a node of T with a child w. Let us call the edge vw of T *necessary* for a tangle leaf $\ell \geq w$ of T if $\beta(vw)$ is a $\leq$-minimal element of the tangle $\lfloor \beta_\ell \rfloor$, or equivalently, of β_ℓ. We call vw *necessary* for a forbidden leaf $\ell \geq w$ if every subset of β_ℓ in $\mathcal{F}$ contains $\beta(vw)$.

Let $(T, r, \beta)_{w \to v} := (T_{w \to v}, r_{w \to v}, \beta_{w \to v})$, where $T_{w \to v}$ is obtained from T by contracting the edge vw of T and deleting any other child w' of v together with the subtree spanned by the nodes $u \geq_r w'$. We continue to use 'w' for the new node constructed from the edge vw, and thus think of $V(T_{w \to v})$ as a subset of $V(T)$. Similarly, we think of $E(T_{w \to v})$ as a subset of $E(T)$ and continue to use 'β' to denote its labelling $\beta_{w \to v}$. If $v = r$ we set $r_{w \to v} := w$; otherwise we keep $r_{w \to v} := r$.

Note that $T_{w \to v}$ inherits the sets E_u from T for its nodes u. In particular, leaves of $T_{w \to v}$ were also leaves of T, and the separations s_u for non-leaf nodes u of $T_{w \to v}$ are still what they were for T. Hence $T_{w \to v}$ is still ordered if T was. However, $T_{w \to v}$ may no longer be thoroughly ordered in $\vec{S}$ even if T was.

The sets β_u remain unchanged for all $u \not\geq w$, while for all $u \geq w$ the sets β_u in $T_{w \to v}$ arise from β_u in T by deleting $\beta(vw)$ from it. (Recall that β was injective on $E(rTu)$, by Definition 3.1.) In particular, $T_{w \to v}$ is consistent if T was. For tangle leaves ℓ of $T_{w \to v}$, the tangle $\lfloor \beta_\ell \rfloor$ of S is the same as it was when β_ℓ was taken in T; its former element $\beta(vw)$ then was a non-minimal element of this tangle.

Let us call a node v of T *necessary in* T if for every child w of v there exists a leaf $\ell \geqslant w$ of T such that vw is necessary for ℓ. A tangle structure tree (T, r, β) is *irreducible* if every node of T is necessary.

Theorem 5.1. *Let $\vec{S}$ be a separation system, and let $\mathcal{F}$ be any set. For every $\mathcal{F}$-tangle structure tree T of $\vec{S}$ there exists an irreducible $\mathcal{F}$-tangle structure tree T' of $\vec{S}$ obtained from T by a sequence $T = T^1, \dots, T^n = T'$ with $T^{i+1} = T^i_{w \to v}$ for suitable nodes v, w of T^i.*

The contraction process in Theorem 5.1 can cause our tangle structure trees from Section 4 to lose their property of being thoroughly ordered. But it preserves the main consequence of that property: that for every leaf ℓ the set β_ℓ is efficient

in $\lfloor \beta_\ell \rfloor$. Let us call separation trees with this property *efficient*. One can show that if the tree T in Theorem 5.1 was thoroughly ordered, then both T and the tree T' contracted from it are efficient.

6. The $\mathcal{F}$-tangle structure theorems

Let us summarize what we have shown so far. Theorem 6.1 below tells us when tangle structure trees exist. It adds 'irreducible' to Theorem 4.2 at the expense of weakening 'thoroughly ordered' to 'efficient' – which, as we shall see in Theorem 6.2, still captures all their key aspects.

Theorem 6.1. *Let $\vec{S}$ be an ordered separation system. If a set $\mathcal{F}$ is standard and rich for $\vec{S}$, then $\vec{S}$ has an efficient and irreducible ordered $\mathcal{F}$-tangle structure tree.*

Theorem 6.2. *Let $\vec{S}$ be an ordered separation system. Every efficient ordered $\mathcal{F}$-tangle structure tree (T, r, β) for $\vec{S}$ has the following properties.*

(i) *For every leaf ℓ of T, the set β_ℓ is consistent in S. If β_ℓ has no subset in $\mathcal{F}$, then $\lfloor \beta_\ell \rfloor$ is a tangle of S, and every $\mathcal{F}$-tangle of S arises in this way.*

(ii) *For every non-leaf node v of T, the set β_v has no subset in $\mathcal{F}$, and s_v is not oriented by $\lfloor \beta_v \rfloor$. Both its orientations $\vec{s} \in \{\overleftarrow{s_v}, \overrightarrow{s_v}\}$ are minimal in $\beta_v \cup \{\vec{s}\}$ in the partial ordering of $\vec{S}$. Its order $|s|$ is maximal in $\{\, |s_u| : u \leq_r v \}$.*

(iii) *For every orientation τ of S there is a unique leaf ℓ of T such that $\beta_\ell \subseteq \tau$. If τ is consistent, then $\lfloor \beta_\ell \rfloor \subseteq \tau$ and β_ℓ is efficient in τ.*

(iv) *For every $\mathcal{F}$-tangle τ of S there is a unique leaf ℓ of T such that $\beta_\ell \subseteq \tau$. Then $\tau = \lfloor \beta_\ell \rfloor$, and β_ℓ is efficient in the tangle τ.*

(v) *For every consistent orientation τ of S that is not a tangle, the set β_ℓ in* (iii) *has a subset σ in $\mathcal{F}$. Every such σ is efficient in τ.*

Let us briefly address the question of how efficient our structure trees are in displaying the $\mathcal{F}$-tangles of S, and the certificates from $\mathcal{F}$ for consistent orientations of S that are not $\mathcal{F}$-tangles.

There is a unique way to encode a single tangle τ efficiently: we have to know its minimal elements, but only these, since the orientations of all the other separations in S are then determined by consistency of τ. The structure tree from Theorem 6.1 does not achieve this for all the tangles of S simultaneously. Indeed, one only needs two separations in S to see that no structure tree can in general display only the minimal elements of all its tangles.

However, since our structure tree T in Theorem 6.1 is irreducible, it does the next best thing. By Theorem 6.2 (iv), the minimal elements of every tangle τ of S are displayed as labels along the path $rT\ell$, where ℓ is the tangle leaf with $\tau = \lfloor \beta_\ell \rfloor$. Conversely, every label $\overrightarrow{s_v}$ of an edge of T is minimal in *some* tangle of S (with a tangle leaf $\ell > v$), or indispensable as an element of the subsets $\sigma \in \mathcal{F}$ of β_ℓ that certify that $\lfloor \beta_\ell \rfloor$ is not a tangle, for some forbidden leaf $\ell > v$.

Finally, let us apply Theorem 6.1 to the special case that S has no $\mathcal{F}$-tangles. If $\mathcal{F}$ consists of stars (of oriented separations; see Section 2), such certificates are known: they are the *S-trees over $\mathcal{F}$* from [14], which generalize the branch-decompositions introduced by Robertson and Seymour [20] for graph tangles.

Our structure trees from Theorem 6.1 provide maximally efficient certificates for the non-existence of $\mathcal{F}$-tangles for general $\mathcal{F}$, not necessarily consisting of stars, when all their leaves are forbidden. Let us call such $\mathcal{F}$-tangle structure trees *$\mathcal{F}$-trees*.

Theorem 6.3. *Let $\vec{S}$ be an ordered separation system, and let $\mathcal{F}$ be any set that is standard and rich for $\vec{S}$. Then exactly one of the following assertions holds:*

(i) *there exists an $\mathcal{F}$-tangle of S;*

(ii) *there exists an $\mathcal{F}$-tree of $\vec{S}$.*

In the case of (ii), the $\mathcal{F}$-tree can be chosen to be irreducible, efficient and ordered.

7. Applications: blocks, profiles, and cluster tangles

Graph tangles can be expressed as $\mathcal{F}$-tangles in the universe of all separations of a graph G: take as $\mathcal{F}$ the set $\mathcal{T}$ of triples $\{\,(A_i, B_i) \mid i = 1, 2, 3\,\}$ of oriented separations such that $G[A_1] \cup G[A_2] \cup G[A_3] = G$. These are not stars. But it is not hard to show [12] that the $\mathcal{T}$-tangles of a graph are precisely its $\mathcal{T}^*$-tangles, where $\mathcal{T}^*$ consists of the sets in $\mathcal{T}$ that are stars.

Due to this history, $\mathcal{F}$-tangles have so far been studied mostly when $\mathcal{F}$ consisted of stars of separations. Such $\mathcal{F}$-tangles already allow for vast generalizations of the original tangles of graphs. However, there are contexts in which $\mathcal{F}$-tangles occur naturally and $\mathcal{F}$ does not consist of stars. Three of these are:

- blocks in graphs, more precisely: k-blocks for all $k \in \mathbb{N}$;
- profiles: the most comprehensive generalization of graph tangles that still admits a tree-of-tangles structure theorem;
- cluster tangles in large data sets.

We shall apply our results to all these in this section, giving a minimum of background for context in each case.

Common to all tangle contexts are two fundamental theorems about tangles in graphs, which both go back to the original paper of Robertson and Seymour [20] in which tangles were first introduced. These have become known [11] as

- the tree-of-tangles theorem; and
- the tangle-tree duality theorem.

The tree-of-tangles theorem, in its simplest form, says that all the maximal tangles in a graph can be distinguished[8] by a nested set T of separations, one that can be represented as the set of separations associated with a tree-decomposition of the graph. The separations in T can be chosen so that they have minimum order for any separation that distinguishes the pair of tangles they distinguish, and *canonically*, which means that T is invariant under all the automorphisms of the graph [11]. These enhanced tree-of-tangle theorems have been further generalized to abstract separation systems [8, 12, 13].

The tangle-tree duality theorem, in its simplest form, says that if a graph has no tangle of some given order k then it has a nested set T of separations of order $< k$ that witnesses the non-existence of those tangles: any orientation τ of the graph's separations of order $< k$ will in particular orient T, and thereby the edges of the decomposition tree associated with T, and following those oriented edges will take us to a node whose incident edges correspond to a set in $\mathcal{T}^*$; thus, τ is not a tangle. The tangle-tree duality theorem, too, has been generalized to abstract separation systems [12, 14]. There are also unifications of both theorems into one [2, 3, 16].

[8] Recall that a separation *distinguishes* two tangles if they orient it differently.

The certificates for the non-existence of an $\mathcal{F}$-tangle offered by the tangle-tree duality theorems require that $\mathcal{F}$ consist of stars. Hence if we look for certificates for the non-existence of $\mathcal{F}$-tangles when $\mathcal{F}$ does not consist of stars, we need a fresh start. Our Theorem 6.3 offers such an alternative.

If $\mathcal{F}$ does consist of stars, one can convert our $\mathcal{F}$-trees of $\vec{S}$ into the traditional certificates for the non-existence of $\mathcal{F}$-tangles for such $\mathcal{F}$, which are *S-trees over $\mathcal{F}$*. Our conversion theorem thus allows us to deduce the most comprehensive known tangle-tree duality theorem [14] from our premise for the existence of tangle structure trees. This differs from the premise of the tangle-tree duality theorem in [14], but follows from it in most scenarios [4, Theorem 7.4], making our result stronger.

For our first new application of Theorem 6.3, let us look at blocks. For any $k \in \mathbb{N}$, a *k-block* in a graph is any maximal set of at least k vertices no two of which can be separated by $< k$ vertices. Every k-block X in a graph G defines an $\mathcal{B}_k$-tangle of S_k for

$$\mathcal{B}_k := \Big\{ \sigma \subseteq \vec{S} : \Big| \bigcap \{ B \mid (A,B) \in \sigma \} \Big| < k \Big\},$$

where S is the set of all separations of G: just orient every separation in S_k towards X. Conversely, given any $\mathcal{B}_k$-tangle τ of S_k in G, the set $\bigcap\{ B \mid (A,B) \in \tau\}$ is a k-block in G, which in turn defines τ as indicated.

One of the earliest applications of what later became $\mathcal{F}$-tangles was a 'tree-of-tangles' theorem for blocks: finite graphs have canonical tree-decompositions that distinguish all their k-blocks [7]. This extended a classical decomposition theorem of Tutte for 2-blocks. Our Theorem 6.3 now also implies a dichotomy theorem:

Theorem 7.1. *Let S be the set of separations of a graph G, with the usual order function [11], and let $k \in \mathbb{N}$. Then exactly one of the following assertions holds:*

(i) *G has a k-block;*

(ii) *there exists an $\mathcal{B}_k$-tree of $\vec{S}_k$.*

In the case of (ii), the $\mathcal{B}_k$-tree can be chosen to be irreducible, efficient, and ordered.

Next, profiles. Profiles of graphs were introduced by Hundertmark [18], who noticed that the above-mentioned 'tree-of-tangles' theorem for blocks [7] required in its proof much less information about blocks than is provided in their definition. All that is needed is that the orientation of S_k which a k-block induces avoids the set $\mathcal{F}$ of triples of the form $\{\vec{r}, \vec{s}, \overleftarrow{r} \vee \overleftarrow{s}\}$. In the context of graphs, this supremum is taken in the universe of all the separations of the given graph.

More generally, given any separation system $\vec{S}$ in some universe $\vec{U}$ of separations, we let

$$\mathcal{P} = \big\{ \{\vec{r}, \vec{s}, \overleftarrow{r} \vee \overleftarrow{s}\} : \vec{r}, \vec{s} \in \vec{U} \big\},$$

and refer to the $\mathcal{P}$-tangles of S as its *profiles*. A profile is *regular* if none of its elements is small.

Hundertmark's discovery was seminal: since graph tangles are profiles too, the canonical construction from [7] for blocks, rewritten for profiles [13], offered canonical trees of tangles – something that Robertson and Seymour [20] had been unable to achieve directly by their tangle methods.

So what about tangle-tree duality for profiles? As in the case of blocks, we cannot have classical tangle-tree duality, because $\mathcal{P}$ does not consist of stars. But Theorem 6.3 now offers another dichotomy theorem, just as for blocks.

Let us call a profile of S *strong* if it has no subset in

$$\mathcal{P}_{\mathrm{s}} = \big\{ \{\vec{r}, \vec{s}, \overleftarrow{t}\} : \vec{r}, \vec{s}, \vec{t} \in \vec{U} \text{ and } \overleftarrow{t} \leq \overleftarrow{r} \vee \overleftarrow{s} \big\}.$$

Strong profiles are regular. And in most contexts they are no stronger than ordinary regular profiles: every regular profile of a submodular separation system in a distributive universe of separations can be shown to be strong. The regular profiles of S, then, are precisely its $\mathcal{P}_s$-tangles.

Theorem 6.3 thus yields the following dichotomy theorem for profiles:

Theorem 7.2. *Let $\vec{S}$ be any submodular separation system in a distibutive universe of separations. Then exactly one of the following assertions holds:*

(i) *there exists a regular profile of S;*

(ii) *there exists an $\mathcal{P}_s$-tree of $\vec{S}$.*

In the case of (ii), the $\mathcal{P}_s$-tree can be chosen to be irreducible, efficient, and ordered with respect to any order function on S.

Finally, let us apply our results to tangles defined by clusters in large data sets. The basic setup is that $\vec{S}$ is the set 2^V of all subsets of our data set V, the involution $*$ given by complementation in V. Then S is the set of bipartitions $s = \{A, B\}$ of V, whose orientations are $\vec{s} = A$ and $\overleftarrow{s} = B$ (say). We then choose an order function on S that assigns high values to separations that divide many pairs of elements of V which we consider as 'close', the idea being that separations of low order cannot divide many close pairs and therefore cannot cut through a cluster, only chip off a few points. This, then, implies that large clusters induce $\mathcal{F}$-tangles for $\mathcal{F}$ defined as

$$\mathcal{C}_n := \big\{ \{ \vec{r}, \vec{s}, \vec{t} \} : |\vec{r} \cap \vec{s} \cap \vec{t}| < n \big\},$$

where the *agreement* value n is chosen to fit the context.

Thus, $\mathcal{C}_n$-tangles in data sets are a bit like $\mathcal{B}_k$-tangles in graphs, except that n is now a value separate from the order k of the tangles considered. An important difference is that, unlike k-blocks in graphs, the clusters captured by those tangles are not normally subsets X of V that we can, or even wish to, specify explicitly: they can be 'fuzzy', and do not have to lie entirely on the side of the separations chosen by the tangle they induce, only mostly. Quantitatively, while the intersection of any three elements of an $\mathcal{C}_n$-tangle contains at least n datapoints, the intersection of all of them will likely be empty – unlike in the case of k-blocks in graphs, which are equal to the intersection of *all* the sides of the graph's separations chosen by the tangle.

In other contexts than generic clustering one sometimes looks for $\mathcal{C}_n$-tangles not of the set S of all the bipartitions of V, but of some specific, hand-designed set S of bipartitions of V. For example, S might be a questionnaire whose elements partition a population V of individuals that have answered it into those that answered yes and those that answered no. In such a context, $\mathcal{C}_n$-tangles can be interpreted as 'typical' ways to answer the questions in S, as *mindsets* found in the population V regarding the questions in S. See [10] for more.

In all these contexts we can now use Theorem 6.1 to efficiently display the clusters or mindsets in V as determined by S, as long as $\mathcal{C}_n$ is standard and rich for $\vec{S}$. If we are interested specifically in certifying that our data set has no clusters of some desired density at all, we can similarly apply Theorem 6.3.

The sets $\mathcal{C}_n$ are standard for all $\vec{S} \neq \emptyset$ and $n > 0$: then $\{\emptyset\} \in \mathcal{C}_n$, and only $\emptyset$ can be co-trivial in $\vec{S}$. Also, $\mathcal{C}_n$ is clearly closed under minimization in $\vec{S}$: replacing the elements $\vec{s}$ of a triple σ with subsets $\vec{s'} \leq \vec{s}$ leads to $\bigcap \sigma' \subseteq \bigcap \sigma$, so $\sigma \in \mathcal{C}_n$ implies $\sigma' \in \mathcal{C}_n$. This easily implies that $\mathcal{C}_n$ is rich for $\vec{S}$.

Theorems 6.1 and 6.3 thus have the following instances in our context:

Corollary 7.3. *Let $\vec{S}$ be any non-empty set of subsets of a set V that is closed under taking complements in V. Let any order function on S be given, and let $n > 0$. Then $\vec{S}$ has an efficient and irreducible ordered $\mathcal{C}_n$-tangle structure tree.* □

Corollary 7.4. *Let $\vec{S}$ be any non-empty set of subsets of a set V that is closed under taking complements in V. Let any order function on S be given, and let $n > 0$. Then exactly one of the following assertions holds:*

(i) *there exists a $\mathcal{C}_n$-tangle of S;*

(ii) *there exists a $\mathcal{C}_n$-tree of $\vec{S}$.*

In the case of (ii), the $\mathcal{C}_n$-tree can be chosen to be irreducible, efficient, and ordered. □

References

[1] *Tangles: e-book and software,* see `tangles-book.com`.

[2] Sandra Albrechtsen, *Refining trees of tangles in abstract separation systems: Inessential parts*, Combinatorial Theory **4(1)** (2024; arXiv:2302.01808).

[3] ———, *Optimal trees of tangles: refining the essential parts*, Innovations in Graph Theory **3** (2026), 171–206.

[4] H.v. Bergen and R. Diestel, *Tangle structure trees II: trees of tangles and tangle-tree duality*, arXiv:2603.17756.

[5] J. Carmesin, R. Diestel, M. Hamann, and F. Hundertmark, *Canonical tree-decompositions of finite graphs I. Existence and algorithms*, J. Comb. Theory Ser. B **116** (2016), 1–24; arXiv:1305.4668.

[6] ———, *Canonical tree-decompositions of finite graphs II. Essential parts*, J. Comb. Theory Ser. B **118** (2016), 268–283; arXiv:1305.4909.

[7] J. Carmesin, R. Diestel, F. Hundertmark, and M. Stein, *Connectivity and tree structure in finite graphs*, Combinatorica **34** (2014), no. 1, 1–35; arXiv:1105.1611.

[8] J. Carmesin and J. Kurkofka, *Entanglements*, J. Comb. Theory, Ser. B **164** (2024), 17–28; arXiv:2205.11488.

[9] R. Diestel, *Abstract separation systems*, Order **35** (2018), 157–170; arXiv:1406.3797.

[10] ———, *Tangles: A structural approach to artificial intelligence in the empirical sciences*, Cambridge University Press, 2024. e-book: `tangles-book.com/book/`.

[11] ———, *Graph theory* (6th edition), Springer-Verlag, 2025.
Electronic edition available at `http://diestel-graph-theory.com/`.

[12] R. Diestel, J. Erde, and D. Weißauer, *Structural submodularity and tangles in abstract separation systems*, J. Comb. Theory Ser. A **167C** (2019), 155–180; arXiv:1805.01439.

[13] R. Diestel, F. Hundertmark, and S. Lemanczyk, *Profiles of separations: in graphs, matroids, and beyond*, Combinatorica **39** (2019), no. 1, 37–75; arXiv:1110.6207.

[14] R. Diestel and S. Oum, *Tangle-tree duality in abstract separation systems*, Advances in Mathematics **377** (2021), 107470; arXiv:1701.02509.

[15] C. Elbracht, J. Kneip, and M. Teegen, *Obtaining trees of tangles from tangle-tree duality*, J. Combinatorics **13** (2022), 251–287; arXiv:2011.09758.

[16] Joshua Erde, *Refining a tree-decomposition which distinguishes tangles*, SIAM J. Discrete Math. **31** (2017), 1529–1551; arXiv:1512.02499.

[17] Martin Grohe and Pascal Schweitzer, *Computing with tangles*, Proceedings of the forty-seventh annual acm symposium on theory of computing, 2015, pp. 683–692.

[18] F. Hundertmark, *Profiles. An algebraic approach to combinatorial connectivity*, unpublished manuscript 2011.

[19] S. Klepper, C. Elbracht, D. Fioravanti, J. Kneip, L. Rendsburg, M. Teegen, and U.v. Luxburg, *Clustering with tangles: Algorithmic framework and theoretical guarantees*, Journal of Machine Learning Research **24** (2023; arXiv:2006.14444), 1–56.

[20] N. Robertson and P.D. Seymour, *Graph minors. X. Obstructions to tree-decomposition*, J. Comb. Theory, Series B **52** (1991), 153–190.

TANGLE STRUCTURE TREES, PART I

(HVB & RD) UNSUBMITTED JOURNAL VERSION

Abstract. We introduce a comprehensive data structure, tangle structure trees, which simultaneously displays all the $\mathcal{F}$-tangles of an abstract separation system for very general obstruction sets $\mathcal{F}$. It simultaneously also displays certificates $\sigma \in \mathcal{F}$ for any non-existence of such tangles, or for the non-extendability of low-order tangles to higher-order ones.

Our theorem can be applied to produce the structures of the classical tree-of-tangles and tangle-tree duality theorems, both for graph tangles and for their known generalizations to more general separation systems. It extends those theorems to obstruction sets $\mathcal{F}$ that need not define profiles (as they must in all known trees of tangles) or consist of stars of separations (as they must in traditional tangle-tree duality).

Our existence proof for these structure trees is constructive. The construction has been implemented in open-source software available for tangle detection and further analysis.

1. Introduction

The notion of 'tangles' was originally introduced by Robertson and Seymour [21] as an abstract concept of high local connectivity in graphs, one that unifies several more concrete such notions, such as highly connected subgraphs or minors. What all these notions have in common is that, given any low-order separation of the graph, any highly cohesive substructure must lie mostly on one of its two sides: since the separation has low order, it cannot split it into two roughly equal halves. A tangle remembers only how all the low-order separations are oriented in this way, each towards that highly cohesive substructure: the collection of all these oriented separations is called a *tangle*. See [11] for a precise definition and basic facts about graph tangles.

Tangles have since been generalized to more general settings than graphs. In the setting of set partitions, they offer a precise theoretical basis for 'fuzzy' real-world problems such as clustering in large datasets [10]. All these, including graphs, are special cases of so-called *abstract separation systems* [9]: an axiomatic setting that assumes only the most basic properties of graph separations. This is the setting in which tangles are most easily and comprehensively treated, and we shall use this framework also in this paper.

There are two fundamental theorems about tangles, which both have their origins in [21] and are also treated in [11]. The *tree-of-tangles* theorem exposes a tree-like structure in the graph or dataset whose tangles we consider, which 'distinguishes' the tangles in that they are shown to 'live in' different areas of this tree-like structure. The other is the *tangle-tree duality* theorem, which certifies the non-existence of a tangle by exposing another tree-like structure in the graph or

Date: October 6, 2026.

dataset, one whose existence clearly precludes the existence of a tangle.[1] This is the more fundamental of the two theorems, in the sense that the other can be reduced to its abstract version (see below) but not vice versa [18] .

Both the above theorems have been generalized to $\mathcal{F}$-tangles of abstract separation systems, ways of simultaneously orienting the 'separations' of a given structure in such a way that no three of those oriented separations form an element of $\mathcal{F}$.[2] This collection $\mathcal{F}$ has to satisfy some constraints that depend on which of the two theorems we are considering.

For the (known) tree-of-tangles theorems, those which display all the $\mathcal{F}$-tangles of a separation system simultaneously, $\mathcal{F}$ has to include triples of oriented separations that are reminiscent of ultrafilters. If our separations are bipartitions of a fixed set V, for example, then $\mathcal{F}$ must contain all triples of the form $\{A, B, \overline{A \cap B}\}$: if A and B are subsets of V deemed 'large', in that the $\mathcal{F}$-tangle orients the partitions $\{A, \overline{A}\}$ and $\{B, \overline{B}\}$ towards them, then the complement of $A \cap B$ cannot also be 'large'. Such $\mathcal{F}$-tangles are called *profiles*, and all existing tree-of-tangle theorems for $\mathcal{F}$-tangles assume that these are profiles [1, 2, 5, 6, 8, 13, 14, 19]. There are even a tree-of-tangles theorems that display profiles of different 'order' simultaneously. Such profiles, however, must be *robust*: they also must not contain certain triples similar to, but slightly different from, those above. More about profiles in Section 7.[3]

For the tangle-tree duality theorems, $\mathcal{F}$ has to satisfy another constraint: its elements have to be *stars*, nested sets of oriented separations pointing towards each other [9]. The tree-like structures by which tangle-tree duality theorems certify the non-existence of $\mathcal{F}$-tangles require that $\mathcal{F}$ consist of stars: if it does not, these structures cannot be tree-like.

This state of the art leaves $\mathcal{F}$-tangles that are not profiles without any known way of displaying all those tangles simultaneously. And if $\mathcal{F}$ does not consist of stars, it leaves separation systems that have no $\mathcal{F}$-tangles without any known way of organizing the elements of $\mathcal{F}$ into a data structure that displays them as easily checkable certificates for the non-existence of such tangles.

The structure trees whose existence we prove in this paper do both these things: they display all the $\mathcal{F}$-tangles of a separation system even when they are not profiles, and for those orientations of the separation system that are not $\mathcal{F}$-tangles they display certificates in $\mathcal{F}$ that show why they are not.[4]

Tangle structure trees display all this information in a single, comprehensive, data structure, which is maximally efficient in the following, structural, sense. A single tangle is most efficiently displayed by listing just those of its elements that are minimal in the poset of oriented separations that comes with a separation system. Indeed, all these are needed to determine the tangle, but any other elements can be deduced from them. Our structure trees display all the $\mathcal{F}$-tangles of a separation system S simultaneously by listing only the oriented separations that are minimal in one of those tangles, including the tangles of subsystems consisting of

[1] In a graph, this would be a tree-decomposition into parts too small for a tangle to live in.

[2] In the case of graph tangles, $\mathcal{F}$ consists of the sets of up to three oriented separations of the graph such that the union of the 'back' sides of these three separations covers the entire graph [11].

[3] In a follow-up paper [4] we show how our structure trees give rise to trees of $\mathcal{F}$-tangles whenever these are robust, even if they are not profiles.

[4] In a follow-up paper [4] we show that if $\mathcal{F}$ does consist of stars, our duality theorem implies the same tangle-tree dichotomy as the classical one.

separations of lower order. And for orientations of S that are not $\mathcal{F}$-tangles it displays certificates from $\mathcal{F}$ inside those same minimal subsets of oriented separations, in a way readily accessible in the structure tree.

Our paper is organized as follows. In Section 2 we provide the necessary background for tangles in abstract separation systems, the framework in which we shall construct our structure trees and prove their existence.

In Section 3 we introduce basic tangle structure trees. In Section 4 we prove their existence, and in Section 5 we show how to make them efficient. In Section 6 we collect all this information together to prove our main results, the '$\mathcal{F}$-tangle structure theorems'. In Theorem 6.5 we note specifically how the main structure theorem provides certificates from $\mathcal{F}$ for the non-existence of $\mathcal{F}$-tangles if there are none.

In Section 7, finally, we apply our results to three particular types of $\mathcal{F}$-tangles. The first two include $\mathcal{F}$-tangles in graphs: those induced by k-blocks, and those that are profiles. In neither of these does the defining set $\mathcal{F}$ consist of stars, so traditional tangle-tree duality fails. We derive new dichotomy theorems for these.

In our third application we treat $\mathcal{F}$-tangles that encode clusters in large datasets. These are neither profiles nor do their obstruction sets $\mathcal{F}$ consist of stars. Our structure trees display all the clusters by way of their $\mathcal{F}$-tangles, while also witnessing the absence of clusters from the other areas of the data set by displaying elements of $\mathcal{F}$.

2. Tangle basics

In this section we give precise definitions and notation for abstract tangles, largely following [9] and indicating any deviations.[5]

Tangles of graphs are ways of orienting their separations, each towards one of its two sides. Abstract tangles are designed to work in scenarios where there need not be anything to 'separate'. In order to retain our intuition from graphs, however, we continue to refer to the things of which our abstract tangles pick one of two variants (which they will indeed do) as 'separations'. These are defined by noting some key properties of graph separations and making them into axioms, as follows.

A *separation system* $(\vec{S}, \leqslant, ^*)$ is a set $\vec{S}$, whose elements we call *oriented separations*, that comes with a partial ordering $\leqslant$ on $\vec{S}$ and an order-reversing involution $^*\colon \vec{S} \to \vec{S}$. Thus, for any two elements[6] $\vec{r}, \vec{s}$ of $\vec{S}$ with $\vec{r} \leqslant \vec{s}$ we have $\vec{r}^* \geqslant \vec{s}^*$. We write $\vec{s}^* =: \overleftarrow{s}$, and call $\overleftarrow{s}$ the *inverse* of $\vec{s}$. While we allow formally that $\vec{s} = \overleftarrow{s}$, in which case we call $\vec{s}$ and s *degenerate*, this does not happen often in practice.[7]

If a separation system $\vec{U}$ happens to be a lattice, that is, if there is a supremum $\vec{r} \vee \vec{s}$ and an infimum $\vec{r} \wedge \vec{s}$ in $\vec{U}$ for every two elements $\vec{r}, \vec{s} \in \vec{U}$, we call $\vec{U}$ a *universe* of separations. It is *distributive* if it is distributive as a lattice. A separation system $\vec{S} \subseteq \vec{U}$ is *submodular* if for every two elements of $\vec{S}$ either their infimum or their supremum in $\vec{U}$ also lies in $\vec{S}$.

Very rarely we may have separations $\vec{s} \leq \overleftarrow{s}$; then $\vec{s}$ is *small* and $\overleftarrow{s}$ is *large*.[8] We say that $\vec{s}$ is *trivial* (and $\overleftarrow{s}$ is *co-trivial*) in $\vec{S}$ if there exists a pair of inverse

[5]The most important difference is that, for historical reasons, the partial ordering on $\vec{S}$ used in [9] is the inverse of ours. So terms like 'large' and 'small', infima and suprema etc, are reversed.

[6]We often denote the elements of $\vec{S}$ by letters with an arrow, in either direction, precisely in order to have a simple way to refer to their dual elements: by reversing the arrow. But the arrow directions have no meaning: an arbitrary element of $\vec{S}$ could be denoted equally as $\vec{s}$ or as $\overleftarrow{s}$.

[7]The only degenerate separation of a graph $G = (V, E)$, for example, is $\{V, V\}$.

[8]The small separations of a graph G are those of the form (V, A) with $A \subseteq V = V(G)$.

separations $\vec{r}, \overleftarrow{r} < \vec{s}$ in $\vec{S}$. Trivial separations are clearly large, so co-trivial ones are small, but the converse need not hold. See [9] for more on these technicalities if desired.

The set of *unoriented separations* in $(\vec{S}, \leq, ^*)$ is

$$S := \{\{\vec{s}, \overleftarrow{s}\} : \vec{s} \in \vec{S}\}.$$

We call the elements $\vec{s}, \overleftarrow{s}$ of s its *orientations*. An *orientation of* S is a set $\tau \subseteq \vec{S}$ that contains exactly one orientation of every $s \in S$. For $s \in S$ we then denote by $\tau(s)$ the unique orientation of s contained in τ. An orientation of a subset of S is a *partial orientation* of S.

If $\vec{r} \geqslant \vec{s}$ we say that $\vec{r}$ *points towards* s (and that $\overleftarrow{r}$ *points away from* s). We say that $\vec{r}$ *points towards* an oriented separation $\vec{s}$ whenever it points towards s, i.e., if $\vec{r} \geqslant \vec{s}$ or $\vec{r} \geqslant \overleftarrow{s}$, and similarly for 'points away from'. A *star* is a set σ of non-degenerate oriented separations that point towards each other. As is easy to check, this happens if and only if $\vec{r} \geqslant \overleftarrow{s}$ (and hence $\vec{s} \geqslant \overleftarrow{r}$) for all distinct $\vec{r}, \vec{s} \in \sigma$.

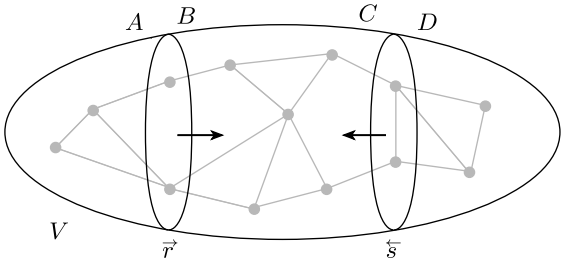


FIGURE 1. Nested separations $r = \{A, B\}$ and $s = \{C, D\}$ of a graph. Their orientations $\vec{r} = (A, B)$ and $\overleftarrow{s} = (D, C)$ point towards each other, since $\vec{r} \geq \vec{s}$ (as $B \supseteq D$) and $\overleftarrow{s} \geq \overleftarrow{r}$ (as $C \supseteq A$).

Two separations $s, r \in S$ are *nested* if they have orientations that are comparable under $\leqslant$. Oriented separations are *nested* if their underlying unoriented separations are nested. A subset of $\vec{S}$ is *nested* if its elements are pairwise nested.

A subset of $\vec{S}$ is *consistent* if no pair of its elements $\vec{r}, \vec{s}$ with $r \neq s$ point away from each other. Stars are examples of consistent nested sets of oriented separations.

We shall often be interested in consistent orientations of S. For each of its elements $\vec{r}$, a consistent orientation of S will also contain every $\vec{s} > \vec{r}$ other than, possibly, $\vec{s} = \overleftarrow{r}$. Consistent partial orientations of S are easily seen to extend to consistent orientations of S, unless they contain a separation that is co-trivial in $\vec{S}$; see [9].

If $\sigma \subseteq \vec{S}$ is consistent, we say that $\vec{s} \in \vec{S}$ is *required* by σ if $\vec{s} \notin \sigma$ and $\sigma \cup \{\overleftarrow{s}\}$ is inconsistent. We shall see in Lemma 2.2 that, pathological cases aside, $\sigma \cup \{\vec{s}\}$ will then be consistent. The *closure* of σ is

$$\lfloor \sigma \rfloor := \sigma \cup \{\vec{s} \in \vec{S} : \vec{s} \text{ is required by } \sigma\}.$$

Note that σ requires $\vec{s} \notin \sigma$ if and only if there exists an $\vec{r} \in \sigma$ such that $r \neq s$ and $\vec{s} > \vec{r}$. Thus

$$\lfloor \sigma \rfloor = \sigma \cup \{\, \vec{s} \in \vec{S} : \exists\, \vec{r} \in \sigma \text{ such that } r \neq s \text{ and } \vec{s} > \vec{r}\},$$

which motivates the term of (upward) 'closure'. If $\lfloor \sigma \rfloor$ is consistent, then $\lfloor \lfloor \sigma \rfloor \rfloor$ is (defined and) easily shown to be equal to $\lfloor \sigma \rfloor$.

If σ contains no small separations, the expression above simplifies to

$$\lfloor \sigma \rfloor = \{ \vec{s} \in \vec{S} : \exists \vec{r} \in \sigma \text{ such that } \vec{r} \leq \vec{s} \}.$$

Indeed, any $\vec{s}$ in this latter set either lies in σ or there is some $\vec{r} \in \sigma$ such that $\vec{r} < \vec{s}$; in that case $r \neq s$, since otherwise $\overleftarrow{s} = \vec{r} < \vec{s}$, making $\vec{r} \in \sigma$ is small.

Lemma 2.1. *Let $\sigma \subseteq \tau \subseteq \vec{S}$ be consistent sets. Then $\lfloor \sigma \rfloor \subseteq \lfloor \tau \rfloor$. If τ is an orientation of all of S, then $\lfloor \tau \rfloor = \tau$.*

Proof. If $\vec{s}$ is required by σ, then $\sigma \cup \{\overleftarrow{s}\}$ is inconsistent. Since $\sigma \subseteq \tau$, it follows that $\tau \cup \{\overleftarrow{s}\}$ is also inconsistent; hence $\vec{s}$ is either an element of or required by τ. Thus every element of $\lfloor \sigma \rfloor$ lies in $\lfloor \tau \rfloor$, proving the first claim.

If τ is an orientation of all of S then $\lfloor \tau \rfloor = \tau$, since for any $\vec{s} \in \lfloor \tau \rfloor \setminus \tau$ the set $\tau \cup \{\overleftarrow{s}\} = \tau$ would be inconsistent. □

Consistent orientations of S can contain small separations. But examples are rare and can be counter-intuitive, so we often exclude them. Co-trivial separations cannot occur in consistent orientations of S. Indeed if $\overleftarrow{s}$ is co-trivial, witnessed by $r \in S$, then every orientation of S will have to orient r too. But it cannot do so consistently with $\overleftarrow{s}$, since both $\vec{r}$ and $\overleftarrow{r}$ are inconsistent with $\overleftarrow{s}$. Similarly, $\lfloor \{\overleftarrow{s}\} \rfloor$ is then inconsistent, since it contains $\vec{r}$ and $\overleftarrow{r}$, which are both inconsistent with $\overleftarrow{s}$.

Lemma 2.2. *Let $\tau \subseteq \vec{S}$ be consistent, and suppose that τ has no elements that are co-trivial in $\vec{S}$. Then $\lfloor \tau \rfloor$ is consistent, and $\lfloor \tau \rfloor \setminus \tau$ contains at most one orientation of any $s \in S$.*

Proof. If $\lfloor \tau \rfloor$ is inconsistent, it contains orientations $\vec{r}, \vec{s}$ of distinct $r, s \in S$ with $\vec{s} < \overleftarrow{r}$. Since τ is consistent, at least one of $\vec{r}, \vec{s}$ is not in τ. Without loss of generality assume $\vec{r} \notin \tau$: Then $\vec{r} \in \lfloor \tau \rfloor \setminus \tau$, so $\vec{r}$ is required by τ (by definition of $\lfloor \tau \rfloor$), which means that $\tau \cup \{\overleftarrow{r}\}$ is inconsistent. Thus there exists $\vec{t} \in \tau$ with $t \neq r$ and $\overleftarrow{r} < \overleftarrow{t}$. We thus have $\vec{s} < \overleftarrow{r} < \overleftarrow{t}$, that is $\vec{s}$ and $\vec{t}$ point away from each other.

If $\vec{s} = \vec{t}$, the above inequality yields $\vec{r}, \overleftarrow{r} < \overleftarrow{t}$, which makes $\vec{t} \in \tau$ co-trivial, contradicting our assumptions. Hence $\vec{s} \neq \vec{t}$ as well as $\vec{s} \neq \overleftarrow{t}$, giving $s \neq t$. The fact that $\vec{s}$ and $\vec{t} \in \tau$ point away from each other thus implies that $\{\vec{s}, \vec{t}\}$ is inconsistent, so $\vec{s} \notin \tau$ as well as $\vec{r} \notin \tau$.

As earlier with $\vec{r}$, the fact that $\vec{s} \in \lfloor \tau \rfloor$ now implies that $\vec{r} < \overleftarrow{s} < \overleftarrow{t'}$ for some $\vec{t'} \in \tau$. The two inequalities together now give $\vec{t'} < \overleftarrow{r} < \overleftarrow{t}$, so $\vec{t'}$ and $\vec{t}$ point away from each other. As τ is consistent and contains both, this means that $t' = t$. As $\vec{t'} < \overleftarrow{t}$, we thus have $\vec{t'} = \vec{t}$. But now $\vec{r} < \overleftarrow{t'} = \overleftarrow{t}$ as well as $\overleftarrow{r} < \overleftarrow{t}$. This makes $\vec{t} \in \tau$ co-trivial in $\vec{S}$, contradicting our assumptions. This proves the first assertion.

For the second, note that τ requires any $\vec{s} \in \lfloor \tau \rfloor \setminus \tau$, by definition of $\lfloor \tau \rfloor$. This means that $\tau \cup \{\overleftarrow{s}\}$ is inconsistent, so $\overleftarrow{s} \notin \lfloor \tau \rfloor$ since $\lfloor \tau \rfloor$ is consistent. □

An *order function* on S is any map $S \to \mathbb{R}$. Unless otherwise mentioned, we denote such order functions as $s \mapsto |s|$. We extend them to $\vec{S}$ by letting $|\vec{s}| := |\overleftarrow{s}| := |s|$. A separation system $(\vec{S}, \leq, ^*)$ given with an order function on S is an *ordered separation system*. For every $k \in \mathbb{R}$ we let

$$\vec{S}_k := \{\vec{s} \in \vec{S} : |s| < k\};$$

this is again an ordered separation system.

We are sometimes interested in orientations of S that do not have certain subsets. We typically collect those together in some set $\mathcal{F}$, whose elements we call *forbidden*

subsets. Formally, if $\mathcal{F}$ is any set, we say that $\tau \subseteq \vec{S}$ *avoids* $\mathcal{F}$ if τ has no subset in $\mathcal{F}$, i.e., if no subset of τ is an element of $\mathcal{F}$.

Definition 2.3. An *$\mathcal{F}$-tangle* of S is an $\mathcal{F}$-avoiding consistent orientation of S. The $\mathcal{F}$-tangles of the subsets S_k of S are the $\mathcal{F}$-tangles *in* $\vec{S}$.

When the context is clear, e.g. when $\mathcal{F}$ is some given set as above, we shall usually say 'tangle' rather than '$\mathcal{F}$-tangle' in this paper. A separation $s \in S$ *distinguishes* two tangles if they orient it differently: if one contains $\vec{s}$, the other $\overleftarrow{s}$.

3. Tangle structure trees

We adopt the graph-theoretic terminology of [11]. A *tree* is a connected acyclic graph. Given nodes t, t' of a tree T we write tTt' for the unique t–t' path in T. A *rooted tree* is a tree with a distinguished node called its *root*. Given a tree T with root r, we define a partial order $\leqslant_r$ on $V(T)$ by declaring $x \leqslant_r y$ if x lies on the path in T from r to y. Maximal elements, including the root if $|T| = 1$, are called *leaves*. Any direct successors in $<_r$ of a node of T are its *children*. We write E_v for the set of edges of T from a node v to its children.

Let $(\vec{S}, \leq, ^*)$ be a separation system, and let $\mathcal{F}$ be any set.

Definition 3.1. A *separation tree* (T, r, β) on $\vec{S}$ consists of a rooted tree (T, r) together with an edge labelling $\beta\colon E(T) \to \vec{S}$ such that for every non-leaf $v \in V(T)$ there exists a separation $s_v \in S$ such that β restricts to a bijection $E_v \to \{\vec{s_v}, \overleftarrow{s_v}\}$ and $s_u \neq s_v$ whenever $u <_r v$.

Thus, every non-leaf node v of a separation tree has either one or two children. If there is no need to refer to r or β explicitly, we usually abbreviate (T, r, β) to T. For every node v, we write β_v for the set $\beta\big(E(rTv)\big)$ of edge labels on the path rTv.

Lemma 3.2. *Let (T, r, β) be a separation tree on $\vec{S}$. Every orientation τ of S contains β_ℓ as a subset for some unique leaf $\ell \in V(T)$.*

Proof. Since $\beta_r = \emptyset \subseteq \tau$, there exists a node $v \in V(T)$ that is maximal with respect to $\leqslant_r$ subject to $\beta_v \subseteq \tau$. If v is not a leaf, let v' be its child with $\beta(vv') = \tau(s_v) =: \vec{s}$. As $\beta_v \subseteq \tau$, and $\vec{s} \in \tau \setminus \beta_v$ since $s_u \neq s_v$ for all $u < v$, we see that $\beta_{v'} = \beta_v \cup \{\vec{s}\} \subseteq \tau$ contradicts the maximality of v. Hence $v =: \ell$ is our desired leaf.

For a proof that ℓ is unique, let ℓ' be another leaf and let v be the greatest common ancestor of ℓ and ℓ'. Then v has distinct children w and w' with $\beta(vw) \in \beta_\ell$ and $\beta(vw') \in \beta_{\ell'}$. Since $\beta(vw) \neq \beta(vw')$, they cannot both lie in τ. As $\beta(vw) \in \beta_\ell \subseteq \tau$ we must have $\beta(vw') \notin \tau$, and hence $\beta_{\ell'} \not\subseteq \tau$. □

A separation tree T is *consistent* if $\beta_v \subseteq \vec{S}$ is consistent for every node $v \in T$.

Lemma 3.3. *Let (T, r, β) be a consistent separation tree on $\vec{S}$, and let v be a non-leaf node of T. Then $\vec{s_v}$ and $\overleftarrow{s_v}$ are $\leq$-minimal in $\beta_v \cup \{\vec{s_v}\}$ and in $\beta_v \cup \{\overleftarrow{s_v}\}$, respectively.*

Proof. Suppose not. Without loss of generality assume there exists $\vec{s} \in \beta_v$ with $\vec{s} < \vec{s_v}$. Let w be the child of v with $\beta(vw) = \overleftarrow{s_v}$. Then $\{\vec{s}, \overleftarrow{s_v}\} \subseteq \beta_w$ is inconsistent, contradicting the consistency of (T, r, β). □

Corollary 3.4. *For every consistent separation tree and non-leaf node v we have $\{\vec{s_v}, \overleftarrow{s_v}\} \cap \lfloor \beta_v \rfloor = \emptyset$. That is, s_v is not oriented by $\lfloor \beta_v \rfloor$.*

Proof. Suppose, say, that $\vec{s_v} \in \lfloor \beta_v \rfloor$. Since $s_u \neq s_v$ for all nodes $u < v$, we know that $\vec{s_v} \notin \beta_v$. Our assumption that $\vec{s_v} \in \lfloor \beta_v \rfloor$ thus means that $\vec{s_v}$ is required by β_v. Hence there exists $\vec{r} \in \beta_v$ such that $\vec{r} < \vec{s_v}$ (and $r \neq s_v$). This contradicts Lemma 3.3. □

Lemma 3.3 and Corollary 3.4 imply that consistent separation trees offer an efficient data structure for storing consistent orientations of S, in particular, for tangles. Indeed, any such orientation τ labels the tree's edges between the root and some leaf ℓ, but this β_ℓ is only a small subset of τ: elements $\vec{s}$ of τ required by β_v for some $v < \ell$ will not appear in $\beta_\ell \setminus \beta_v$. But they can be reconstructed from β_v by the consistency of τ.

By definition of separation trees, the separations labelling their edges do not repeat along any root-to-leaf path: if $u < v$ then $s_u \neq s_v$. Hence the length of any path from the root to a leaf is at most $|S|$. Since every node has at most two children, this bounds the size of the tree:

Corollary 3.5. *Separation trees on $\vec{S}$ have at most $2^{|S|}$ leaves and fewer than $2^{|S|+1}$ nodes.* □

A leaf ℓ of a consistent separation tree on $\vec{S}$ is an *$\mathcal{F}$-tangle leaf*, or *tangle leaf* for short, if the closure $\lfloor \beta_\ell \rfloor$ of β_ℓ is an $\mathcal{F}$-tangle of S. A non-leaf node v is a *tangle node* if there is a tangle leaf $\ell \geq_r w$ for every successor w of v in $>_r$.

A leaf ℓ is *forbidden* (by $\mathcal{F}$) if β_ℓ contains an element of $\mathcal{F}$ as a subset. Note that tangle leaves are never forbidden.

Definition 3.6. An *$\mathcal{F}$-tangle structure tree of $\vec{S}$* is a consistent separation tree on $\vec{S}$ in which every leaf is either a tangle leaf or forbidden, and for every non-leaf node v the set β_v has no subset in $\mathcal{F}$.

Every $\mathcal{F}$-tangle structure tree of $\vec{S}$ displays all the $\mathcal{F}$-tangles of S, not just some of them:

Theorem 3.7. *Let $\vec{S}$ be a separation system, let $\mathcal{F}$ be any set, and let T be any $\mathcal{F}$-tangle structure tree of $\vec{S}$.*

(i) *For every $\mathcal{F}$-tangle τ of S there is a unique leaf ℓ of T such that $\lfloor \beta_\ell \rfloor = \tau$.*
(ii) *If all the leaves of T are forbidden, then S has no $\mathcal{F}$-tangle.*

Proof. (i) By Lemma 3.2 there is a unique leaf ℓ with $\beta_\ell \subseteq \tau$. This leaf is not forbidden, since $\beta_\ell \subseteq \tau$ has no subset in $\mathcal{F}$. As T is a tangle structure tree, this means that ℓ is a tangle leaf: that $\lfloor \beta_\ell \rfloor$ is a tangle of S. As $\lfloor \beta_\ell \rfloor \subseteq \tau$, this tangle can only be τ.

(ii) is immediate from (i). □

In the next section we shall determine the sets $\mathcal{F}$ for which S admits a tangle structure tree whose tree-order is compatible with a given order function on S, in that lower-order separations label edges further down on the tree. As we shall see in Theorem 6.3 and Corollary 6.4, such trees display not only all the $\mathcal{F}$-tangles of the entire set S, as all tangle structure trees do by Theorem 3.7, but all the $\mathcal{F}$-tangles in $\vec{S}$: the $\mathcal{F}$-tangles of the sets $S_k = \{s \in S : |s| < k\}$ with $k \in \mathbb{N}$.

4. Existence of tangle structure trees

Let $\vec{S}$ be an ordered separation system, and let $\mathcal{F}$ be any set. We shall need two conditions on $\mathcal{F}$ to ensure that $\vec{S}$ has a tangle structure tree.

The first is that $\{\overleftarrow{s}\} \in \mathcal{F}$ for every $\vec{s} \in \vec{S}$ that is trivial in $\vec{S}$. This holds for all sets $\mathcal{F}$ of interest, and if it does we call $\mathcal{F}$ *standard for* $\vec{S}$. Intuitively, co-trivial separations $\overleftarrow{s}$ point to places in our graph or structure that are tiny – too small to house any tangle. More formally, we already saw in Section 2 that no consistent orientation of S can contain co-trivial separations. Assuming that $\mathcal{F}$ is standard for $\vec{S}$ therefore places no restrictions on the $\mathcal{F}$-tangles of S.

The second condition has more substance. If T is any separation tree on $\vec{S}$ then, by Lemma 3.2, every orientation τ of S contains β_ℓ for a unique leaf ℓ. If τ is consistent and $\lfloor\beta_\ell\rfloor$ orients all of S, then $\lfloor\beta_\ell\rfloor = \tau$ by Lemma 2.1. If τ is not a tangle, then $\tau = \lfloor\beta_\ell\rfloor$ has a subset σ in $\mathcal{F}$ that witnesses this. If T is in fact a tangle structure tree, we know that our consistent orientation τ of S is a tangle (without having to assume that $\lfloor\beta_\ell\rfloor$ orients all of S) unless it contains such a set $\sigma \in \mathcal{F}$ not only in $\lfloor\beta_\ell\rfloor$ but even in β_ℓ: among the edge labels of T.

In order for this to be possible, we therefore need to make some 'richness' assumption about $\mathcal{F}$: an assumption which, in our example, ensures that $\mathcal{F}$ has enough elements to contain a subset also of β_ℓ as soon as it contains a subset of $\lfloor\beta_\ell\rfloor$.

We shall identify an essentially weakest-possible such richness condition below, in Definition 4.3. This will need some more preparation. For motivation, readers are invited to peek at Lemma 4.4 and the definition preceding it now. That definition clearly implies that $\mathcal{F}$ contains a subset of β_ℓ whenever it contains a subset of $\lfloor\beta_\ell\rfloor$. The notion of richness from Definition 4.3 will still imply this, but is weaker.

A separation tree T on an ordered separation system $\vec{S}$ is *ordered* if $|s_v| \leq |s_w|$ whenever v and w are non-leaves of T with $v \leq w$. It is *thoroughly ordered* (in $\vec{S}$) if, for every non-leaf node v, the separation s_v is not oriented by $\lfloor\beta_v\rfloor$[9] and has minimum order among the separations in S not oriented by $\lfloor\beta_v\rfloor$.

Lemma 4.1. *Every thoroughly ordered separation tree is ordered.*

Proof. Let T be a thoroughly ordered separation tree on $\vec{S}$. If it is not ordered, it has nodes v, w with $v < w$ such that $|s_v| > |s_w|$. Since s_w is not oriented by $\lfloor\beta_v\rfloor \subseteq \lfloor\beta_w\rfloor$, this contradicts the requirement that s_v have minimum order among the separations not oriented by $\lfloor\beta_v\rfloor$. □

We say that $\vec{s} \in \vec{S}$ is *weakly eclipsed* by $\vec{r} \in \vec{S}$ if $\vec{r} < \vec{s}$ and $|r| \leq |s|$, and *eclipsed* by $\vec{r}$ if $\vec{r} < \vec{s}$ and $|r| < |s|$. Given any set $\tau \subseteq \vec{S}$, a subset $\sigma \subseteq \tau$ is *efficient* (in τ) if no element of σ is eclipsed by any other element of τ. It is *strongly efficient* if no element of σ is weakly eclipsed by any other element of τ. Note that if the order function on S is injective and τ is a partial orientation of S, then every efficient subset of τ is strongly efficient in τ.[10]

Lemma 4.2. *Let v be a node of a thoroughly ordered separation tree T on $\vec{S}$. Then*

(i) *every strongly efficient subset of $\lfloor\beta_v\rfloor$ is contained in β_v;*

(ii) *if β_v is consistent, it is efficient in any partial orientation τ of S that includes $\lfloor\beta_v\rfloor$.*

[9] If T is consistent, then this holds by Corollary 3.4.

[10] We need the assumption on τ, since $\vec{s}$ eclipses $\overleftarrow{s}$ weakly if $\vec{s} < \overleftarrow{s}$. So we do not want $\vec{s}, \overleftarrow{s} \in \tau$.

Proof. (i) Let σ be any strongly efficient subset of $\lfloor\beta_v\rfloor$, and let $\vec{s} \in \sigma$ be given. Suppose that $\vec{s} \notin \beta_v$. Then $\vec{s} \in \lfloor\beta_v\rfloor \setminus \beta_v$ is required by β_v, so there exists $\vec{r} \in \beta_v$ with $\vec{r} < \vec{s}$ (and $r \neq s$). As $\vec{r} \in \beta_v$, we have $r = s_u$ for some $u < v$. Choose such an $\vec{r}$ with u minimal in $<_r$.

Our aim is to show that $|r| \leq |s|$: then $\vec{r} \in \lfloor\beta_v\rfloor$ eclipses $\vec{s} \in \sigma$ weakly, contradicting the strong efficiency of σ as a subset of $\lfloor\beta_v\rfloor$. Since T is thoroughly ordered, we shall have $|r| \leq |s|$ as desired if $\lfloor\beta_u\rfloor$ does not orient s: this would make s a candidate for s_u, so $|r| > |s|$ would contradict the fact that $r = s_u$.

So let us show that neither $\vec{s}$ nor $\overleftarrow{s}$ lies in $\lfloor\beta_u\rfloor$. For $\vec{s} \notin \lfloor\beta_u\rfloor$, recall first that $\vec{s} \notin \beta_v \supseteq \beta_u$. Hence if $\vec{s} \in \lfloor\beta_u\rfloor$, there exists an $\vec{r'} \in \beta_u$ such that $\vec{r'} < \vec{s}$. This satisfies $r' = s_{u'}$ for some $u' <_r u$, so $\vec{r'} < \vec{s}$ contradicts our original choice of $\vec{r}$ given $\vec{s}$. Hence $\vec{s} \notin \lfloor\beta_u\rfloor$ as desired. Suppose now that $\overleftarrow{s} \in \lfloor\beta_u\rfloor$. As $\vec{s} > \vec{r}$, any $\vec{r'} \leqslant \overleftarrow{s}$ in β_u satisfies $\vec{r'} \leqslant \overleftarrow{s} < \overleftarrow{r}$. So r was already oriented by $\lfloor\beta_u\rfloor$, contradicting the fact that $r = s_u$. Hence neither $\vec{s}$ nor $\overleftarrow{s}$ lies in $\lfloor\beta_u\rfloor$, as desired.

(ii) Suppose not; then some $\vec{r} \in \tau$ eclipses some $\vec{s} \in \beta_v$. Let $u < v$ be such that $s = s_u$. As $|r| < |s|$ and T is thoroughly ordered, we know that $\lfloor\beta_u\rfloor$ orients r.

Since τ is a partial orientation of S containing $\vec{r}$, it cannot also contain $\overleftarrow{r}$. We thus cannot have $\overleftarrow{r} \in \lfloor\beta_u\rfloor$, since $\lfloor\beta_u\rfloor \subseteq \lfloor\beta_v\rfloor \subseteq \tau$ by Lemma 2.1. So $\vec{r} \in \lfloor\beta_u\rfloor$. As $\vec{r} < \vec{s}$, this implies $\vec{s} \in \lfloor\beta_u\rfloor$ unless $\vec{r} = \overleftarrow{s} \in \beta_u$. Both these contradict the fact that $s = s_u$. □

A typical application of Lemma 4.2 (ii) is that β_v is efficient in any tangle of S that contains it. Such a tangle will be consistent and hence include $\lfloor\beta_v\rfloor$ by Lemma 2.1.

A typical application of Lemma 4.2 (i) will be that witnesses $\sigma \in \mathcal{F}$ to the fact that some $\lfloor\beta_v\rfloor$ fails to extend to a tangle can be found not only in $\lfloor\beta_v\rfloor$ but in β_v itself. Then our structure trees will display such witnesses $\sigma \in \mathcal{F}$ in the label sets β_ℓ of their forbidden leaves ℓ. For this to work with the help of Lemma 4.2, we need $\mathcal{F}$ to contain witnesses that are strongly efficient in their $\lfloor\beta_v\rfloor$.

This motivates our formal richness condition on $\mathcal{F}$:

Definition 4.3. A set $\mathcal{F}$ is *rich for* $\vec{S}$ if every consistent orientation of S that has a subset in $\mathcal{F}$ also has a strongly efficient[11] subset in $\mathcal{F}$.

The assumption that $\mathcal{F}$ is rich will be central to our results. It is needed to ensure that our $\mathcal{F}$-tangle structure trees exist, and we shall see after Lemma 4.7 that no weaker condition on $\mathcal{F}$ will ensure the same.

Although it may look a bit technical at first glance, this 'richness' requirement on $\mathcal{F}$ is quite natural, given the role of these $\mathcal{F}$ in tangle theory. For example, given a consistent orientation τ of S that has a subset $\sigma \in \mathcal{F}$, we can often obtain a strongly efficient subset $\sigma' \in \mathcal{F}$ of τ simply by replacing every element $\vec{s}$ of σ by some $\vec{s'} \leq \vec{s}$ that is minimal in τ. One still has to check then that σ' is indeed in $\mathcal{F}$. But as the idea behind those forbidden triples in $\mathcal{F}$ is that they identify areas in our graph or other structure that are 'too small to be home to a tangle', it is not unnatural for this particular σ' to be in $\mathcal{F}$ if σ was.

Let us cast this example in the form of a lemma. Let us call a set $\mathcal{F}$ *closed under minimization* in a subset τ of $\vec{S}$ if it contains every set $\sigma' \subseteq \tau$ obtained from some $\sigma \subseteq \tau$ in $\mathcal{F}$ by replacing every $\vec{s} \in \sigma$ with some $\vec{s'} \leq \vec{s}$.

[11] in this orientation of S

Lemma 4.4. *If $\mathcal{F}$ is closed under minimization in every consistent orientation of S, then $\mathcal{F}$ is rich for $\vec{S}$.*

Proof. Let τ be any consistent orientation of S that has a subset σ in $\mathcal{F}$. We have to find a set $\sigma' \subseteq \tau$ in $\mathcal{F}$ that is strongly efficient in τ.

Let σ' be obtained from σ by replacing every $\vec{s} \in \sigma$ by some $\vec{s'} \leq \vec{s}$ that is minimal in τ. The set σ' is strongly efficient in τ, since any $\vec{r} < \vec{s'} \in \sigma'$ in τ contradicts the minimal choice of $\vec{s'} \leq \vec{s}$. Since $\mathcal{F}$ is closed under minimization in τ, we have $\sigma' \in \mathcal{F}$ as required. □

When we apply Lemma 4.4 later in Section 7, we shall in fact prove that the sets $\mathcal{F}$ needed there are closed under minimization in all of $\vec{S}$. Note that this is stronger than being closed under minimization in every consistent orientation of S.

Lemma 4.5. *Let $\mathcal{F}$ be rich and standard for $\vec{S}$. Let T be a thoroughly ordered separation tree on $\vec{S}$, and let $v \in V(T)$. Assume that β_v is consistent and avoids $\mathcal{F}$, and that $\lfloor \beta_v \rfloor$ orients all of S. Then $\lfloor \beta_v \rfloor$ is a tangle of S.*

Proof. As β_v avoids $\mathcal{F}$, which is standard, β_v has no element that is co-trivial in $\vec{S}$. As β_v is consistent, $\lfloor \beta_v \rfloor$ is consistent by Lemma 2.2. Since $\lfloor \beta_v \rfloor$ orients all of S, it can thus only fail to be a tangle of S if it has a subset σ in $\mathcal{F}$. As $\mathcal{F}$ is rich for $\vec{S}$, we can choose σ to be strongly efficient as a subset of $\lfloor \beta_v \rfloor$.[12] Then even $\sigma \subseteq \beta_v$ by Lemma 4.2 (i), which contradicts our assumption that β_v avoids $\mathcal{F}$. Hence $\lfloor \beta_v \rfloor$ is a tangle of S. □

Theorem 4.6. *Let $\vec{S}$ be an ordered separation system, and let $\mathcal{F}$ be a set that is rich and standard for $\vec{S}$. Then there exists a thoroughly ordered $\mathcal{F}$-tangle structure tree of $\vec{S}$.*

Proof. Start with the one-node tree $T_0 = \{r\}$ and let $\beta^0 = \emptyset$. We iteratively build trees $T_0 \subsetneq T_1 \subsetneq \dots$ with maps β^i so that each (T_i, r, β^i) is consistent and thoroughly ordered, and none of their sets β^i_v with v a non-leaf node will have a subset in $\mathcal{F}$. Note that T_0 has all these properties. The last of those trees will be our desired tangle structure tree.

If for some n the tree (T_n, r, β^n) is already a tangle structure tree, we are done. Otherwise T_n has a leaf v which is neither a tangle leaf nor forbidden. By Lemma 4.5, $\lfloor \beta^n_v \rfloor$ does not orient all of S; let $s \in S$ be a separation of minimum order not oriented by $\lfloor \beta^n_v \rfloor$.

Form T_{n+1} by adding two children v_1, v_2 at v. Let β^{n+1} agree with β^n on the edges of T_n, and pick orientations $\vec{s} =: \beta^{n+1}(vv_1)$ and $\overleftarrow{s} =: \beta^{n+1}(vv_2)$ of s; then $s = s_v$ in T_{n+1}. By construction, $(T_{n+1}, r, \beta^{n+1})$ is a thoroughly ordered separation tree. It is consistent, because T_n was and neither orientation of s lies in $\lfloor \beta^{n+1}_v \rfloor$. Its only non-leaf node v that was not already a non-leaf node in T_n is v. Since v was not forbidden as a leaf of T_n, the set $\beta^{n+1}_v = \beta^n_v$ has no subset in $\mathcal{F}$.

This process strictly increases $|V(T_n)|$ at each step, but by Corollary 3.5 there is an upper bound on the size of any consistent separation tree in terms of $|S|$. Hence the process terminates after finitely many steps with a thoroughly ordered tangle structure tree. □

[12] If $\lfloor \beta_v \rfloor$ contains both orientations of some $s \in S$, delete one of them to obtain a consistent orientation of S as required in the definition of 'rich' for $\mathcal{F}$.

If our order function on S is injective, Theorem 4.6 has a converse, which shows that our requirement of richness for $\mathcal{F}$ is weakest possible to ensure the existence of a tangle structure tree. This is established by our next lemma:

Lemma 4.7. *Let $\vec{S}$ be a separation system with an injective order function on S, and let $\mathcal{F}$ be any set. If there exists a thoroughly ordered $\mathcal{F}$-tangle structure tree T of $\vec{S}$, then $\mathcal{F}$ is rich for $\vec{S}$.*

Proof. Let τ be any consistent orientation of $\vec{S}$ that has a subset in $\mathcal{F}$. We shall find an efficient subset $\sigma \in \mathcal{F}$ of τ, which will even be strongly efficient since our order function is injective. By Lemma 3.2, T has a leaf ℓ with $\beta_\ell \subseteq \tau$. In particular, β_ℓ is consistent, so $\lfloor \beta_\ell \rfloor \subseteq \tau$ by Lemma 2.1.

Since T is a tangle structure tree, ℓ is either a tangle leaf or forbidden. If ℓ is a tangle leaf then $\lfloor \beta_\ell \rfloor$ is a tangle of all of S. But then $\lfloor \beta_\ell \rfloor = \tau$, which contradicts the fact that τ has a subset in $\mathcal{F}$.

Thus ℓ is forbidden, so β_ℓ has a subset σ in $\mathcal{F}$. By Lemma 4.2 (ii) this σ is efficient in τ, as desired. □

Note that if our order function on S is not injective, the proof of Lemma 4.7 still goes through as stated except for one aspect: the efficient subset $\sigma \in \mathcal{F}$ of τ it finds may not be strongly efficient in τ (as our definition of 'rich' requires).

If our order function on S is injective, the sets $\mathcal{F}$ for which $\vec{S}$ admits a thoroughly ordered tangle structure tree are thus precisely the rich ones:

Theorem 4.8. *Let $\vec{S}$ be a separation system with an injective order function on S, and let $\mathcal{F}$ be a set that is standard for $\vec{S}$. There exists a thoroughly ordered $\mathcal{F}$-tangle structure tree of $\vec{S}$ if and only if $\mathcal{F}$ is rich for $\vec{S}$. If one exists, it is unique.*

Proof. Existence is immediate from Theorem 4.6 and Lemma 4.7. Uniqueness follows from the definition of 'thoroughly ordered' and our assumption that the order function on S is injective. □

Let us illustrate the above by an example. Consider a separation system whose elements are the non-empty subsets of a 4-element set V, with $\leq$ defined as $\subseteq$, and * as complementation in V. Let r, s be two crossing separations; their orientations are 2-sets. The 1-element subsets of V then are precisely those of the form $\vec{r} \cap \vec{s}$ for suitable orientations $\vec{r}, \vec{s}$ of r and s. We call the four unoriented separations which each partition V into a singleton and a 3-set the *corners* of r and s. Let S be the set of r, s, and their four corners.

The set S then has four *principal* consistent orientations, those that orient one corner t away from r and s and all the other elements of S towards t. And it has the *non-principal* consistent orientations, which orient r and s arbitrarily and all four corners towards r and s. Choose an injective order function on S such that $|r| < |s| < |t|$ for all corners t.

Let us first consider as $\mathcal{F}$ the 4-element set $\mathcal{P} = \big\{ \{ \vec{r}, \vec{s}, \overleftarrow{t} \} \mid \vec{t} = \vec{r} \cap \vec{s} \big\}$.[13] The unique thoroughly ordered separation tree T on $\vec{S}$ then starts with orienting r and s, as a tree T' with four leaves, one for every orientation of $\{r, s\}$. For each of these leaves v, the next separation s_v to be oriented is one of the corners t. If $\vec{t} = \vec{r} \cap \vec{s}$, say, we have $\vec{t} = \beta(v\ell)$ for a tangle leaf ℓ for the principal tangle $\lfloor \beta_\ell \rfloor = \lfloor \vec{t} \rfloor$. And

[13] We shall study these $\mathcal{P}$-tangles more closely in Section 7.

$\overleftarrow{t} = \beta(v\ell')$ for a forbidden leaf ℓ', with $\beta_{\ell'} = \{\vec{r}, \vec{s}, \overleftarrow{t}\} \in \mathcal{P}$. As this happens at each of the four leaves v of T', our T is a tangle structure tree of $\vec{S}$ for $\mathcal{F} = \mathcal{P}$.

Now let $\mathcal{F}$ be the set $\mathcal{R}$ of all triples in $\vec{S}$ whose complements partition V.[14] These are the triples of the form either $\{\vec{r}, \overleftarrow{t}, \overleftarrow{t'}\}$ with $\vec{t} = \vec{r} \cap \vec{s}$ and $\vec{t'} = \vec{r} \cap \overleftarrow{s}$, or $\{\vec{s}, \overleftarrow{t}, \overleftarrow{t'}\}$ with $\vec{t} = \vec{s} \cap \vec{r}$ and $\vec{t'} = \vec{s} \cap \overleftarrow{r}$. Our thoroughly ordered separation tree T then starts with the same T' as before. Every leaf v of T' once more sends an edge $v\ell$ to a tangle leaf ℓ with $\beta(v\ell) = \vec{t}$ and $\mathcal{R}$-tangle $\lfloor \beta_\ell \rfloor = \lfloor \vec{t} \rfloor$. But with $\beta_v = \{\vec{r}, \vec{s}\}$ as earlier we now have $\beta_{\ell'} = \{\vec{r}, \vec{s}, \overleftarrow{t}\} \notin \mathcal{R}$ for the other child ℓ' of v. But note that $\overleftarrow{t'} \supset \vec{s}$ for the corner $\vec{t'} = \vec{r} \cap \overleftarrow{s}$. Hence $\lfloor \beta_{\ell'} \rfloor$ is an orientation of S with a subset in $\mathcal{R}$: the set $\{\vec{r}, \overleftarrow{t}, \overleftarrow{t'}\}$. Thus, ℓ' is neither a tangle leaf nor forbidden, and T is not an $\mathcal{R}$-tangle structure tree.

As predicted by Theorem 4.8, this difference between $\mathcal{P}$ and $\mathcal{R}$ is reflected by the fact that $\mathcal{P}$ is rich for $\vec{S}$ but $\mathcal{R}$ is not: the consistent orientation $\lfloor \beta_{\ell'} \rfloor$ of S has a subset in $\mathcal{R}$, but it has no efficient subset in $\mathcal{R}$. In particular, its subset $\{\vec{r}, \overleftarrow{t}, \overleftarrow{t'}\} \in \mathcal{R}$, which prevented T from being an $\mathcal{R}$-tangle structure tree, is not efficient, because $\vec{s} \subset \overleftarrow{t'}$ eclipses $\overleftarrow{t'}$.

In Section 5 we shall prove that our tangle structure trees can be improved further, without any additional requirements on $\mathcal{F}$, by contracting 'inessential' edges: edges whose label $\vec{s}$ is neither needed in any sets $\sigma \in \mathcal{F}$ witnessing that a leaf is forbidden, nor needed to determine any tangles.

Unfortunately, this contraction process can cause our structure trees to lose their property of being thoroughly ordered. We shall therefore extract from this property its essence, a slightly weaker property to be called 'efficiency', which is still strong enough to make the branches $rT\ell$ and their label sets β_ℓ as efficient for displaying the tangles $\lfloor \beta_\ell \rfloor$ as they are in thoroughly ordered structure trees, but weak enough to survive the contraction process we envisage for Section 5.

Given an ordered separation system $\vec{S}$, we call a separation tree (T, r, β) *efficient* if for every leaf ℓ the set β_ℓ is efficient in $\lfloor \beta_\ell \rfloor$.

Lemma 4.9. *Let $\vec{S}$ be an ordered separation system. Then every thoroughly ordered separation tree (T, r_0, β) on $\vec{S}$ is efficient.*

Proof. Consider any leaf ℓ of T. If β_ℓ is not efficient in $\lfloor \beta_\ell \rfloor$ as claimed, then some $\vec{s} \in \beta_\ell$ is eclipsed by some $\vec{r} \in \lfloor \beta_\ell \rfloor$. Let $v, w < \ell$ be such that $\vec{s} = \vec{s_v}$ and $\vec{r} \geq \vec{s_w} \in \beta_\ell$.[15]

As $\vec{r}$ eclipses $\vec{s}$, we have $\vec{r} < \vec{s}$ and $|r| < |s|$; in particular, $r \neq s = s_v$. The fact that T is thoroughly ordered thus implies that r is already oriented by $\lfloor \beta_v \rfloor$. We cannot have $\vec{r} \in \lfloor \beta_v \rfloor$, since this would place $\vec{s_v} = \vec{s} > \vec{r}$ or $\overleftarrow{s_v}$ in $\lfloor \beta_v \rfloor$ too,[16] which would contradict the fact that T is thoroughly ordered. Thus, $\overleftarrow{r} \in \lfloor \beta_v \rfloor$; let $u < v$ be such that $\overleftarrow{r} \geq \vec{s_u} \in \beta_v$.

If $u = w$, then $\vec{s_u} = \vec{s_w} \leq \vec{r} < \vec{s}$ and hence $\vec{s_v} = \vec{s} \in \lfloor \beta_v \rfloor$ or $\overleftarrow{s_v} \in \beta_v$,[16] both of which contradict the fact that T is thoroughly ordered.

If $u < w$, then $\overleftarrow{s_w} \geq \overleftarrow{r} \geq \vec{s_u} \in \beta_w$. Then $\overleftarrow{s_w} \in \lfloor \beta_w \rfloor$ unless $\vec{s_w} \in \beta_w$,[16] which both contradict the fact that T is thoroughly ordered.

[14] We shall study these $\mathcal{R}$-tangles more closely in [4].

[15] Recall that our arrow notation for oriented separations is never fixed. We are thus free to use forward arrows to denote the orientations of s_v and s_w that lie in β_ℓ.

[16] By the definition of $\lfloor\ \rfloor$, the assumption of $\vec{s} > \vec{r} \in \lfloor \beta_v \rfloor$ implies $\vec{s} \in \lfloor \beta_v \rfloor$ only if $\overleftarrow{s} \notin \beta_v$.

If $w < u$, finally, then $\overleftarrow{s_u} \geq \vec{r} \geq \vec{s_w} \in \beta_u$, so $\overleftarrow{s_u} \in \lfloor \beta_u \rfloor$ unless $\vec{s_u} \in \beta_u$,[16] which both contradict the fact that T is thoroughly ordered. □

5. Irreducible and efficient tangle structure trees

Let (T, r, β) be a consistent separation tree on a separation system $\vec{S}$. Let v be a node of T with a child w.

Let us call the edge vw of T *necessary* for a tangle leaf $\ell \geq w$ of T if $\beta(vw)$ is a $\leq$-minimal element of the tangle $\lfloor \beta_\ell \rfloor$, or equivalently, of β_ℓ. We call vw *necessary* for a forbidden leaf $\ell \geq w$ if every subset of β_ℓ in $\mathcal{F}$ contains $\beta(vw)$.

Let $(T, r, \beta)_{w \to v} := (T_{w \to v}, r_{w \to v}, \beta_{w \to v})$, where $T_{w \to v}$ is obtained from T by contracting the edge vw of T and deleting any other child w' of v together with the subtree spanned by the nodes $u \geq_r w'$. We continue to use 'w' for the new node constructed from the edge vw, and thus think of $V(T_{w \to v})$ as a subset of $V(T)$. Similarly, we think of $E(T_{w \to v})$ as a subset of $E(T)$ and continue to use 'β' to denote its labelling $\beta_{w \to v}$. If $v = r$ we set $r_{w \to v} := w$; otherwise we keep $r_{w \to v} := r$.

Note that $T_{w \to v}$ inherits the sets E_u from T for its nodes u. In particular, leaves of $T_{w \to v}$ were also leaves of T, and the separations s_u for non-leaf nodes u of $T_{w \to v}$ are still what they were for T. Hence $T_{w \to v}$ is still ordered if T was. However, $T_{w \to v}$ may no longer be thoroughly ordered in $\vec{S}$ even if T was.

The sets β_u remain unchanged for all $u \not\geq w$, while for all $u \geq w$ the sets β_u in $T_{w \to v}$ arise from β_u in T by deleting $\beta(vw)$ from it. (Recall that β was injective on $E(rTu)$, by Definition 3.1.) In particular, $T_{w \to v}$ is consistent if T was. For tangle leaves ℓ of $T_{w \to v}$, the tangle $\lfloor \beta_\ell \rfloor$ of S is the same as it was when β_ℓ was taken in T; its former element $\beta(vw)$ then was a non-minimal element of this tangle.

Lemma 5.1. *Let $\vec{S}$ be a separation system, and let $\mathcal{F}$ be any set. Let (T, r_0, β) be a consistent separation tree on $\vec{S}$. Let v be a node of T with a successor w, let $\vec{s} := \beta(vw)$, and let $\ell \geq w$ be a leaf. Then vw is necessary for ℓ if and only if $\beta_\ell \setminus \{\vec{s}\}$ has no subset in $\mathcal{F}$ and $\lfloor \beta_\ell \setminus \{\vec{s}\} \rfloor$ is not a tangle of S.*

Proof. Suppose first that vw is necessary for ℓ, and that ℓ is a tangle leaf. Then $\lfloor \beta_\ell \rfloor$ is a tangle of S, so $\beta_\ell \setminus \{\vec{s}\}$ has no subset in $\mathcal{F}$.

For a proof that $\lfloor \beta_\ell \setminus \{\vec{s}\} \rfloor$ is not a tangle of S, we show that it contains neither $\vec{s}$ nor $\overleftarrow{s}$. To see $\vec{s} \notin \lfloor \beta_\ell \setminus \{\vec{s}\} \rfloor$, note that any $\vec{r} \in \beta_\ell \setminus \{\vec{s}\}$ with $\vec{r} \leq \vec{s}$ would satisfy $\vec{r} < \vec{s}$, since also $\vec{s} \in \beta_\ell$ and hence $r \neq s$ by Definition 3.1. Then $\vec{s}$ would not be minimal in β_ℓ, which it is since vw is necessary for the tangle leaf ℓ. And $\overleftarrow{s} \notin \lfloor \beta_\ell \setminus \{\vec{s}\} \rfloor$, since $\lfloor \beta_\ell \rfloor$ is a tangle of S containing $\vec{s}$. Thus, $\lfloor \beta_\ell \setminus \{\vec{s}\} \rfloor$ does not orient s, so it is not a tangle of S.

Suppose, second, that vw is necessary for ℓ, and that ℓ is a forbidden leaf. Then $\beta_\ell \setminus \{\vec{s}\}$ has no subset in $\mathcal{F}$. Suppose $\lfloor \beta_\ell \setminus \{\vec{s}\} \rfloor$ is a tangle of S. This tangle cannot contain $\vec{s}$: it would then contain the entire set β_ℓ, but this has a subset in $\mathcal{F}$ since ℓ is a forbidden leaf. But neither can $\lfloor \beta_\ell \setminus \{\vec{s}\} \rfloor$ contain $\overleftarrow{s}$: then $\beta_\ell \setminus \{\vec{s}\}$ would contain some $\vec{r} \leqslant \overleftarrow{s}$, which would contradict the consistency of β_ℓ, since $r \neq s$ by Definition 3.1 and $\vec{s} \in \beta_\ell$ would thus point away from $\vec{r} \in \beta_\ell$.

Suppose, third, that vw is not necessary for ℓ and that ℓ is a tangle leaf. Then $\lfloor \beta_\ell \rfloor$ is a tangle and $\vec{s}$ is not a minimal element of β_ℓ. We complete our proof in this case by showing that $\lfloor \beta_\ell \setminus \{\vec{s}\} \rfloor = \lfloor \beta_\ell \rfloor$. If $\vec{s} \notin \beta_\ell$ this holds trivially. If $\vec{s}$ lies in β_ℓ but is not minimal in it, then there exists some $\vec{r} \in \beta_\ell$ with $\vec{r} < \vec{s}$. This

implies $\vec{s} \in \lfloor \beta_\ell \setminus \{\vec{s}\} \rfloor$, and hence

$$\lfloor \beta_\ell \rfloor = \lfloor \beta_\ell \setminus \{\vec{s}\} \rfloor \cup \lfloor \{\vec{s}\} \rfloor \subseteq \lfloor \lfloor \beta_\ell \setminus \{\vec{s}\} \rfloor \rfloor = \lfloor \beta_\ell \setminus \{\vec{s}\} \rfloor$$

as desired, unless $\overleftarrow{s} \in \lfloor \beta_\ell \setminus \{\vec{s}\} \rfloor$. But in that case the tangle $\lfloor \beta_\ell \rfloor$ contains both $\vec{s}$ and $\overleftarrow{s}$, which it cannot.

Suppose finally that vw is not necessary for ℓ and that ℓ is a forbidden leaf. Then $\beta_\ell \setminus \{\vec{s}\}$ has a subset in $\mathcal{F}$, as desired. □

Lemma 5.2. *Let $\vec{S}$ be a separation system, and let $\mathcal{F}$ be any set. Let (T, r, β) be a tangle structure tree of $\vec{S}$. Let v be a node of T with a successor w. Then $(T, r, \beta)_{w \to v}$ is a tangle structure tree of $\vec{S}$ if and only if vw is not necessary for any leaf $\ell \geq w$ of T.*

Proof. Suppose first that vw is necessary for some leaf $\ell \geq w$ of T. This ℓ is a leaf also of $T_{w \to v}$. By Lemma 5.1, however, ℓ is neither a tangle leaf nor a forbidden leaf of $T_{w \to v}$. Hence $T_{w \to v}$ is not a tangle structure tree of $\vec{S}$.

Conversely, suppose that vw is not necessary for any leaf $\ell \geqslant w$ in T. These ℓ are precisely (all) the leaves of $T_{w \to v}$. Since β_ℓ taken in $T_{w \to v}$ is obtained from β_ℓ taken in T by deleting $\beta(vw)$, these ℓ are either forbidden or tangle leaves of $T_{w \to v}$ by Lemma 5.1. Thus, $T_{w \to v}$ is a tangle structure tree of $\vec{S}$. □

Let us note a surprising consequence of Lemma 5.2 and its proof. Consider a node v of T with children w, w', and assume that vw is not necessary for any leaf $\ell \geq w$ of T. The truth of this assertion is determined entirely by the values under β of the subtree of T that consists of the path rTv, the subtree spanned by all nodes $u \geq w$, and the edge vw. In particular, it is not affected by the values under β of the edges of the subtree of T spanned by the nodes $u \geq w'$: the subtree we delete when we form $T_{w \to v}$ from T.

By Lemma 5.2, this $T_{w \to v}$ is a tangle structure tree. Every tangle of S thus equals $\lfloor \beta_\ell \rfloor$ for some leaf ℓ of $T_{w \to v}$, which we recall is the same tangle of S that was associated with ℓ in T. (In particular, $\lfloor \beta_\ell \rfloor$ is the same in $T_{w \to v}$ as in T.) So what about the tangles of S whose associated leaf of T lay in the deleted subtree, tangles that include $\beta_{w'}$ and in particular the inverse $\beta(vw')$ of $\beta(vw)$?

The surprising conclusion is that such tangles cannot exist: that if vw is not necessary for any leaf $\ell \geq w$, then $\beta_{w'}$ does not extend to a tangle of S.[17] This is true regardless of whether β_w extends to a tangle of S or not. But in particular, v cannot have been a tangle leaf of T; indeed, T had no tangle node $v' \geq w'$:

Corollary 5.3. *Tangle structure trees T and $T_{w \to v}$ as in Lemma 5.2 have the same sets of tangle nodes.*

It is not hard to construct examples of tangle structure trees that have unnecessary edges. Indeed in many naturally occurring tangle structure trees most edges are unnecessary, and applying Theorem 5.5 below reduces them substantially.

Let us call a node v of T *necessary in T* if for every child w of v there exists a leaf $\ell \geqslant w$ of T such that vw is necessary for ℓ.

Lemma 5.4. *Let $\vec{S}$ be a separation system, and let $\mathcal{F}$ be any set. Let (T, r, β) be a tangle structure tree of $\vec{S}$, and let v be a node of T. Then v has a child w such that $(T, r, \beta)_{w \to v}$ is a tangle structure tree of $\vec{S}$ if and only if v is not necessary in T.*

[17] This does not mean that $\beta(vw')$ cannot lie in a tangle of S: such tangles just do not include β_v.

Proof. Suppose first that v is necessary in T. Let w be any child of v in T. Since v is necessary, there is a leaf $\ell \geq w$ such that vw is necessary for ℓ. By Lemma 5.2, $T_{w \to v}$ is not a tangle structure tree of $\vec{S}$.

Conversely, suppose v is not necessary in T. Then it has a child w such that vw is not necessary for any leaf $\ell \geqslant w$ in T. By Lemma 5.2, $T_{w \to v}$ is a tangle structure tree of $\vec{S}$. □

A tangle structure tree (T, r, β) is called *irreducible* if every node of T is necessary. The following theorem summarizes how tangle structure trees can be contracted and pruned to irreducible trees that inherit many of their properties. The statement of the theorem assumes our notational convention that $V(T_{w \to v}) \subseteq V(T)$.

Theorem 5.5. *Let $\vec{S}$ be a separation system, and let $\mathcal{F}$ be any set. For every $\mathcal{F}$-tangle structure tree T of $\vec{S}$ there exists an irreducible $\mathcal{F}$-tangle structure tree T' of $\vec{S}$ obtained from T by a sequence $T = T^1, \dots, T^n = T'$ with $T^{i+1} = T^i_{w \to v}$ for suitable nodes v, w of T^i.*

The tree T' imposes the same partial ordering on its node as T does. In particular, T' is ordered if T is. Its leaves ℓ are also leaves of T, and their incident edges in T are still edges of T'.

The tangle leaves ℓ of T' are precisely those of T. Their associated tangles $\lfloor \beta_\ell \rfloor$ are the same in T' as in T. The tangle nodes of T' are precisely those of T. A tangle node v separates two tangle leaves in T' if and only if it separates them in T.

Proof. If T^i is not irreducible, pick a node $v \in T^i$ that is not necessary. By Lemma 5.4 it has a child w for which $T^i_{w \to v}$ is a tangle structure tree. Let $T^{i+1} := T^i_{w \to v}$. For the second paragraph of the theorem, note that edges of the form $v\ell$, with ℓ a leaf, are necessary in every tangle structure tree: by Corollary 3.4 if ℓ is a tangle leaf, and by Definition 3.6 if ℓ is forbidden.

For the third paragraph, recall our discussion following Lemma 5.2 and iterate Corollary 5.3. Recall that the separation s_v associated with a tangle node v is the same in T' as in T, and that s_v distinguishes two tangles if and only if v separates their tangle leaves in any tangle structure tree of $\vec{S}$. □

The irreducible tangle structure tree T' obtained in the proof of Theorem 5.5 is not unique. This is because we can check the nodes v of T for suppression in any order, and this order will affect their necessity in the tree T^i in which they are considered for suppression. But it is easy to make T' unique. For example, it is not hard to show that if we consider the nodes v in any linear order that extends their tree order $<_r$ from T – for example, level by level – the resulting tree T' will be the same.

If $\mathcal{F}$ is standard, the $\mathcal{F}$-tangle structure trees we considered in Section 4 were all efficient, because they were thoroughly ordered (Lemma 4.9). The irreducible structure trees we obtained from them in Theorem 5.5 may no longer be thoroughly ordered: as noted before, this property can get lost in the contraction process. But the efficiency of the original structure trees is maintained:

Lemma 5.6. *If T' is obtained from an efficient tangle structure tree T as in Theorem 5.5, then T' too is efficient.*

Proof. Consider $T^{i+1} = T^i_{w \to v}$ as in Theorem 5.5 for $i = 1, \dots, n-1$ in turn. Assuming that T^i is efficient, we have to show that so is T^{i+1}. Write $\beta^{i+1} := \beta^i_{w \to v}$,

where β^i is the edge labelling of T^i. Given any leaf ℓ of T^{i+1}, we have to show that β^{i+1}_ℓ is efficient in $\lfloor \beta^{i+1}_\ell \rfloor$. This follows from the fact that ℓ is also a leaf of T^i, the assumed efficiency of T^i, and the fact that $\beta^{i+1}_\ell \subseteq \beta^i_\ell$. □

6. The $\mathcal{F}$-tangle structure theorems

In this section we summarize what we have shown so far, in a few, concise, statements that we think of as the main results of this paper.

The first of these tells us when $\mathcal{F}$-tangle structure trees exist:

Theorem 6.1. *Let $\vec{S}$ be an ordered separation system. If a set $\mathcal{F}$ is standard and rich for $\vec{S}$, then $\vec{S}$ has an efficient and irreducible ordered $\mathcal{F}$-tangle structure tree.*

Proof. By Theorem 4.6, $\vec{S}$ has a thoroughly ordered tangle structure tree (T, r, β). By Lemma 4.9 it is efficient. The desired structure tree can be obtained from (T, r, β) by Theorem 5.5; it is efficient by Lemma 5.6. □

The tangle structure trees constructed for the proof of Theorem 6.1 have a host of properties designed to help with the tangle analysis of a given separation system. It might have been natural to list these properties as part of Theorem 6.1. However, since the latter is an existence theorem, this would have allowed us to establish these properties just for the structure tree we constructed in its proof.

We would like to stress the fact that those properties are common to all efficient and irreducible ordered tangle structure trees: these terms were designed to encode precisely those properties. Let us show this next:

Theorem 6.2. *Let $\vec{S}$ be an ordered separation system. Every efficient ordered $\mathcal{F}$-tangle structure tree (T, r, β) for $\vec{S}$ has the following properties.*

(i) *For every leaf ℓ of T, the set β_ℓ is consistent in S. If β_ℓ has no subset in $\mathcal{F}$, then $\lfloor \beta_\ell \rfloor$ is a tangle of S, and every $\mathcal{F}$-tangle of S arises in this way.*

(ii) *For every non-leaf node v of T, the set β_v has no subset in $\mathcal{F}$, and s_v is not oriented by $\lfloor \beta_v \rfloor$. Both its orientations $\vec{s} \in \{\overleftarrow{s}_v, \overrightarrow{s}_v\}$ are minimal in $\beta_v \cup \{\vec{s}\}$ in the partial ordering of $\vec{S}$. Its order $|s|$ is maximal in $\{\, |s_u| : u \leq_r v \}$.*

(iii) *For every orientation τ of S there is a unique leaf ℓ of T such that $\beta_\ell \subseteq \tau$. If τ is consistent, then $\lfloor \beta_\ell \rfloor \subseteq \tau$ and β_ℓ is efficient in τ.*

(iv) *For every $\mathcal{F}$-tangle τ of S there is a unique leaf ℓ of T such that $\beta_\ell \subseteq \tau$. Then $\tau = \lfloor \beta_\ell \rfloor$, and β_ℓ is efficient in the tangle τ.*

(v) *For every consistent orientation τ of S that is not a tangle, the set β_ℓ in* (iii) *has a subset σ in $\mathcal{F}$. Every such σ is efficient in τ.*

Proof. (i) The first two statements are immediate from Definition 3.6: tangle structure trees are consistent, and leaves that are not forbidden are tangle leaves. The third assertion is part of (iv) and will be proved there.

(ii) The first assertion is again part of Definition 3.6. The second is Corollary 3.4. The minimality assertion is Lemma 3.3. The maximality assertion holds, because (T, r, β) is ordered.

(iii) By Lemma 3.2, T has a unique leaf ℓ such that $\beta_\ell \subseteq \tau$. If τ is consistent, then $\lfloor \beta_\ell \rfloor \subseteq \tau$ by Lemma 2.1. The claim now follows by Lemma 4.2 (ii).

(iv) We have to show that the leaf ℓ from (iii) satisfies the inclusion $\lfloor \beta_\ell \rfloor \subseteq \tau$ provided there with equality. As τ has no subset in $\mathcal{F}$, the leaf ℓ is not forbidden.

It is therefore a tangle leaf, which means that $\lfloor \beta_\ell \rfloor$ is a tangle of S. Since distinct tangles of S cannot contain one another, this tangle can only be τ.

(v) Consider the leaf ℓ provided for τ by (iii). If τ is not a tangle, then neither is $\lfloor \beta_\ell \rfloor \subseteq \tau$: if it was, then $\lfloor \beta_\ell \rfloor$ and τ would both be orientations of S, implying $\lfloor \beta_\ell \rfloor = \tau$ with a contradiction. Hence β_ℓ has a subset in $\mathcal{F}$, by (i). The efficiency claim follows from (iii), since all subsets of β_ℓ are efficient in τ if β_ℓ is. □

Let us briefly address the question of how efficient our structure trees are in displaying the $\mathcal{F}$-tangles of S, as well as certificates from $\mathcal{F}$ for consistent orientations of S that are not $\mathcal{F}$-tangles.

There is a unique way to encode a single tangle τ efficiently: we have to know its minimal elements, but only these, since the orientations of all the other separations in S are then determined by consistency of τ.[18] The structure tree from Theorem 6.1 does not achieve this for all the tangles of S simultaneously. Indeed, one only needs two separations in S to see that no structure tree can in general display only the minimal elements of all its tangles.[19]

However, since our structure tree T in Theorem 6.1 is irreducible, it does the next best thing. By Theorem 6.2 (iv), the minimal elements of every tangle τ of S are displayed as labels along the path $rT\ell$, where ℓ is the tangle leaf with $\tau = \lfloor \beta_\ell \rfloor$. Conversely, every label $\vec{s_v}$ of an edge of T is minimal in *some* tangle of S (with a tangle leaf $\ell > v$), or indispensable as an element of the subsets $\sigma \in \mathcal{F}$ of β_ℓ that certify that $\lfloor \beta_\ell \rfloor$ is not a tangle, for some forbidden leaf $\ell > v$.

Theorem 6.1 is best possible in another sense too. Its characterizing condition for the existence of tangle structure trees is, essentially, that $\mathcal{F}$ must be rich. This cannot be weakened: if our order function on S is injective, which in real-world applications is the rule rather than the exception, then by Lemma 4.7 the existence of any ordered tangle structure tree (efficient and irreducible or not) implies that $\mathcal{F}$ is rich. The richness condition on $\mathcal{F}$, thus, captures precisely what we need for a tangle structure tree to exist.

If $\mathcal{F}$ is rich not only for $\vec{S}$ but also for some (or all) $\vec{S_k}$, Theorem 6.1 applied in $\vec{S_k}$ says that there exist efficient and irreducible ordered tangle structures trees of all these $\vec{S_k}$. But we do not have to obtain these by independent applications of Theorem 6.1, separately for each k: we can find them all in the original tangle structure tree T for $\vec{S}$ provided by Theorem 4.6, before we applied Theorem 5.5 to 'reduce' it. Let us now see how.

Given an ordered separation tree (T, r, β) on $\vec{S}$ and $k \in \mathbb{R}$, let

$$(T, r, \beta)|_k = (T|_k, r, \beta|_k)$$

[18] Apply Lemma 2.2 to the set σ of its minimal elements to reobtain the entire tangle as $\tau = \lfloor \sigma \rfloor$.

[19] For example, let S consist of just two nested separations, r and s. Let $\mathcal{F} = \emptyset$, so that all three consistent orientations of S are $\mathcal{F}$-tangles. The orientations $\vec{r}$ of r and $\overleftarrow{s}$ of s that point to each other then form a tangle in which they are both minimal, so they must both label an edge of T. Now if we flip either one of them, we obtain a tangle in which the other is still an element, but no longer minimal. If $|r| < |s|$, say, then $\{\vec{r}, \vec{s}\}$ would be such a tangle which our structure tree displays by using both $\vec{r}$ and $\vec{s}$ as labels of a path from the root to a tangle leaf, although $\vec{s}$ alone already determines this tangle. The tangle $\{\overleftarrow{r}, \overleftarrow{s}\}$ would be displayed by the same tree more efficiently, with $\overleftarrow{r}$ labelling an edge from the root to a leaf ℓ, with $\beta_\ell = \{\overleftarrow{r}\}$ and $\lfloor \beta_\ell \rfloor = \{\overleftarrow{r}, \overleftarrow{s}\}$.

be defined as follows. Since T is ordered, its edges e with $|\beta(e)| < k$ form a subtree of T rooted at r, which we take as $T|_k$. As its edge labelling $\beta|_k$ we take the restriction of β to these edges.

Theorem 6.3. *Let (T, r, β) be a thoroughly ordered $\mathcal{F}$-tangle structure tree of $\vec{S}$. Let $k \in \mathbb{R}$, and assume that $\mathcal{F}$ is rich for $\vec{S}_k$. Then $(T, r, \beta)|_k$ is a thoroughly ordered $\mathcal{F}$-tangle structure tree of $\vec{S}_k$.*

Proof. It is clear from the definition of $T|_k$ that it inherits from T the properties of being a thoroughly ordered and consistent separation tree on $\vec{S}_k$. Moreover, non-leaf nodes v of $T|_k$ are non-leaf nodes also of T, so β_v has no subset in $\mathcal{F}$.

For a proof that $T|_k$ is a tangle structure tree it remains to show that every leaf v of $T|_k$ is either a tangle leaf or a forbidden leaf of $T|_k$. If it is neither, then by Lemma 4.5 applied to the closure $\lfloor \beta_v \rfloor_k$ of β_v in $\vec{S}_k$ does not orient all of S_k.

Such a node v cannot be a leaf of T. Indeed, since β_v has no subset in $\mathcal{F}$, it would be a tangle leaf of T, so $\lfloor \beta_v \rfloor$ would be a tangle of S, and $\lfloor \beta_v \rfloor \cap \vec{S}_k = \lfloor \beta_v \rfloor_k$ would be a tangle of S_k (contrary to our assumption). So there exists an $s_v \in S$ whose orientations label the edges e and e' from v to its children in T.

Since T is thoroughly ordered and there exists a separation in S_k not oriented by $\lfloor \beta_v \rfloor_k = \lfloor \beta_v \rfloor \cap \vec{S}_k$, we know that $|s_v| < k$. But this means that e and e' are edges of $T|_k$, contradicting the fact that v is a leaf of $T|_k$. □

Note that the structure trees $T|_k$ found inside T by Theorem 6.3 are nested: for $i < j$ we have $T|_i \subseteq T|_j$, indeed $T|_i = (T|_j)|_i$. This hierarchy of structure trees commutes with the following well-known hierarchy of tangles. For all integers $i < j$, every tangle τ_j of S_j induces a tangle τ_i of S_i, as $\tau_i = \tau_j \cap \vec{S}_i$. A tangle of any S_k that is not induced by a tangle of any $S_{k'} \supsetneq S_k$ is *maximal* in $\vec{S}$. Every tangle in $\vec{S}$, say τ_k of order k, now corresponds to a tangle leaf $\ell(\tau_k)$ of $T|_k$ in that $\tau_k = \lfloor \beta_{\ell(\tau_k)} \rfloor_k$. All these $\ell(\tau_k)$ are nodes of T.

Conversely, if our order function on S is injective, then *every* non-leaf node v of T is such a tangle leaf $\ell(\tau_k)$ of $T|_k$ for some tangle τ_k in $\vec{S}$, since β_v has no subset in $\mathcal{F}$ by Theorem 6.2 (ii): simply choose k just big enough that $\beta(uv) < k$ for u the parent of v. The positions of these nodes in T then commute with the relationship of their tangles: if τ_j induces τ_i for $i \leq j$, then $\ell(\tau_i) \leq_r \ell(\tau_j)$ in the tree-order of T.

Thus, all the maximal tangles in $\vec{S}$ are displayed by the same structure tree:

Corollary 6.4. *Let (T, r, β) be a thoroughly ordered $\mathcal{F}$-tangle structure tree of $\vec{S}$. Assume that the order function on S is injective, and that $\mathcal{F}$ is rich for every $\vec{S}_k$.*

Let T' be obtained from T by deleting any pairs ℓ, ℓ' of forbidden leaves of T that are children of the same node. Then T' is a thoroughly ordered consistent separation tree on $\vec{S}$, which displays precisely the maximal $\mathcal{F}$-tangles in $\vec{S}$ as $\lfloor \beta_\ell \rfloor_k$ for leaves ℓ of T' and suitable k. □

See [4, Section 6] for more details on T' and applications of Corollary 6.4.

We remark that it is not enough in Theorem 6.3 and Corollary 6.4 to assume that $\mathcal{F}$ is rich for $\vec{S}$, rather than for all the relevant $\vec{S}_k$. To see this, recall the example with $\mathcal{F} = \mathcal{R}$ discussed after Theorem 4.8. Let $k \in \mathbb{R}$ be large enough that $S = S_k$ for the separation system $\vec{S}$ discussed there. The unique thoroughly ordered separation tree T constructed there was not an $\mathcal{F}$-tangle structure tree, because it had leaves that were neither forbidden nor tangle leaves. Correspondingly, $\mathcal{F}$ was not rich for this $\vec{S} = \vec{S}_k$.

Now extend this separation system by adding a new separation x of order k, making both its orientations incomparable with all the elements of $\vec{S}$, and adding them to $\mathcal{F}$ as singleton sets $\{\vec{x}\}$ and $\{\overleftarrow{x}\}$. This makes the extended $\mathcal{F}'$ rich for the extended $\vec{S'}$, because $\{\vec{x}\}$ and $\{\overleftarrow{x}\}$ are efficient and every orientation of S' contains one of them. The unique $\mathcal{F}'$-tangle structure tree of $\vec{S'}$ is obtained from T by adding two new leaves at every leaf of T, labelling their incident edges with $\vec{x}$ and $\overleftarrow{x}$. Note that all these leaves are forbidden.

Deleting them, however, as in Corollary 6.4, returns the original T. Contrary to what the corollary claims, T fails to display the maximal $\mathcal{F}'$-tangles in $\vec{S'}$ as $\lfloor\beta_\ell\rfloor$ for its leaves ℓ, since these are the $\mathcal{F}$-tangles of S.

In the proof of Theorem 6.3 we made use of the assumption that T is thoroughly ordered. This holds for the tangle structure trees provided by Theorem 4.6, and these are already efficient by Lemma 4.9. But while their efficiency is maintained when we apply Theorem 5.5 to 'reduce' them, as in the proof of Theorem 6.1, the property of being thoroughly ordered is lost. For the trees $T|_k$ in Theorem 6.3 to be tangle structure trees, however, it is essential that T is thoroughly ordered.[20]

But the tangle structure trees of $\vec{S}_k$ that Theorem 6.3 finds in T can be reduced afterwards: we simply apply Theorem 5.5 to those $T|_k$. This yields irreducible, efficient, ordered tangle structure trees for all the $\vec{S}_k$, the same that Theorem 6.1 would give us if we applied it directly in $\vec{S}_k$.

Similarly, we can mimic the reduction process established in Section 5 to contract the separation tree T' from Corollary 6.4, which displays all the maximal tangles in $\vec{S}$ as $\lfloor\beta_\ell\rfloor$ for suitable leaves ℓ, to an efficient and 'irreducible' such tree. All it needs to facilitate this is that we adapt the definition of 'tangle leaves', and those of ensuing terms such as 'necessary' and 'irreducible', to refer to the maximal tangles in $\vec{S}$, rather than just the tangles of all of S.[21]

Finally, let us apply Theorem 6.1 to the special case that S has no $\mathcal{F}$-tangles. In this case it is desirable to be able to certify this efficiently.

In the special case that $\mathcal{F}$ consists of stars (of oriented separations; see Section 2), such certificates are known: they are the *S-trees over $\mathcal{F}$* from [16], which generalize the branch-decompositions introduced by Robertson and Seymour [21] for graph tangles to $\mathcal{F}$-tangles in abstract separation systems.

Our structure trees from Theorem 6.1 provide maximally efficient certificates for the non-existence of $\mathcal{F}$-tangles for general $\mathcal{F}$, not necessarily consisting of stars, when all their leaves are forbidden. Let us call such $\mathcal{F}$-tangle structure trees *$\mathcal{F}$-trees*.

Theorem 6.5. *Let $\vec{S}$ be an ordered separation system, and let $\mathcal{F}$ be any set that is standard and rich for $\vec{S}$. Then exactly one of the following assertions holds:*

(i) *there exists an $\mathcal{F}$-tangle of S;*

(ii) *there exists an $\mathcal{F}$-tree of $\vec{S}$.*

[20] In the reduction process we contracted or deleted edges whose label was 'unnecessary' for T to display the $\mathcal{F}$-tangles of S and any certificates in $\mathcal{F}$ of their non-existence. But such edges may have been necessary for $T|_k$ to display the tangles of S_k. As an extreme example, take $\mathcal{F} = \emptyset$ and assume that all the oriented separations that are $\leq$-minimal in tangles of S have order k. While S may have elements of lower order that can be oriented to form lower-order tangles, those tangles will not be visible in T, since $\beta(E(T)) \cap \vec{S}_k = \emptyset$.

[21] The only (small) addition in substance needed here occurs early in the proof of Lemma 5.1, where we now have to show also that $\lfloor\beta_\ell \setminus \{\vec{s}\}\rfloor$ cannot be a maximal tangle in $\vec{S}$ other than τ. But this is clear, since $\lfloor\beta_\ell \setminus \{\vec{s}\}\rfloor \subseteq \tau$.

In the case of (ii)*, the $\mathcal{F}$-tree can be chosen to be irreducible, efficient and ordered.*

Proof. By Theorem 3.7 (ii), any $\mathcal{F}$-tree of S precludes the existence of an $\mathcal{F}$-tangle of $\vec{S}$, so we cannot have both (i) and (ii).

Let us prove that (i) or (ii) holds. By Theorem 6.1, $\vec{S}$ has an efficient and irreducible ordered $\mathcal{F}$-tangle structure tree T. If (i) fails, then T has no tangle leaf. Then all its leaves are forbidden, so T is an $\mathcal{F}$-tree. □

Let us close this section with an interesting fact exhibited by $\mathcal{F}$-trees. If S has no tangle, then orientations of S cannot fail to be a tangle just because they are inconsistent: even inconsistent orientations of S must have a subset in $\mathcal{F}$. This is clearly not the case when S does have tangles, e.g. if $\mathcal{F} = \emptyset$, as soon it contains two nested separations (which we can orient away from each other).

Corollary 6.6. *Let $\vec{S}$ be an ordered separation system, and let $\mathcal{F}$ be any set that is standard and rich for $\vec{S}$. If S has no $\mathcal{F}$-tangle, then every orientation of S has a subset in $\mathcal{F}$.*

Proof. Let τ be any orientation of S. By Theorem 6.5, S has an $\mathcal{F}$-tree (T, r, β). Let $rT\ell$ be the maximal path in T from the root such that $\beta(uv) \in \tau$ for its adjacent vertices $u < v$. Since ℓ is a forbidden leaf, $\mathcal{F}$ has an element $\sigma \subseteq \beta_\ell \subseteq \tau$. □

7. Applications: blocks, profiles, and cluster tangles

Graph tangles can be expressed as $\mathcal{F}$-tangles in the universe of all separations of a graph G: take as $\mathcal{F}$ the set $\mathcal{T}$ of triples $\{ (A_i, B_i) \mid i = 1, 2, 3 \}$ of oriented separations such that $G[A_1] \cup G[A_2] \cup G[A_3] = G$. These are not stars. But it is not hard to show [13] that the $\mathcal{T}$-tangles of a graph are precisely its $\mathcal{T}^*$-tangles, where $\mathcal{T}^*$ consists of the sets in $\mathcal{T}$ that are stars.

Due to this history, $\mathcal{F}$-tangles have so far been studied mostly when $\mathcal{F}$ consisted of stars of separations. Such $\mathcal{F}$-tangles already allow for vast generalizations of the original tangles of graphs. However, there are contexts in which $\mathcal{F}$-tangles occur naturally and $\mathcal{F}$ does not consist of stars. Three of these are:

- blocks in graphs, more precisely: k-blocks for any $k \in \mathbb{N}$;
- profiles: the most comprehensive generalization of graph tangles that still admits a tree-of-tangles structure theorem;
- cluster tangles in large datasets.

We shall apply our results to all these in this section, giving a minimum of background for context in each case.

Common to all tangle contexts are two fundamental theorems about tangles in graphs, which both go back to the original paper of Robertson and Seymour [21] in which tangles were first introduced. These have become known [11] as

- the tree-of-tangles theorem; and
- the tangle-tree duality theorem.

The tree-of-tangles theorem, in its simplest form, says that all the maximal tangles in a graph can be distinguished[22] by a nested set T of separations, one that can be represented as the set of separations associated with a tree-decomposition

[22]Recall that a separation *distinguishes* two tangles if they orient it differently.

of the graph. The separations in T can be chosen so that they have minimum order for any separation that distinguishes the pair of tangles they distinguish, and *canonically*, which means that T is invariant under all the automorphisms of the graph [11]. These enhanced tree-of-tangle theorems have been further generalized to abstract separation systems [8, 13, 14].

The tangle-tree duality theorem, in its simplest form, says that if a graph has no tangle of some given order k then it has a nested set T of separations of order $< k$ that witnesses the non-existence of those tangles: any orientation τ of the graph's separations of order $< k$ will in particular orient T, and thereby the edges of the decomposition tree associated with T, and following those oriented edges will take us to a node whose incident edges correspond to a set in $\mathcal{T}^*$; thus, τ is not a tangle. The tangle-tree duality theorem, too, has been generalized to abstract separation systems [13, 16]. There are also unifications of both theorems into one [1, 2, 19].

The certificates for the non-existence of an $\mathcal{F}$-tangle offered by the tangle-tree duality theorems require that $\mathcal{F}$ consist of stars. Indeed, since T is nested, the edges at any node of the decomposition tree that represents T will map to a nested set of separations too, and orienting them towards that node will yield a star of separations. Hence if we look for certificates for the non-existence of $\mathcal{F}$-tangles when $\mathcal{F}$ does not consist of stars, we need a fresh start. Our Theorem 6.5 offers such an alternative.

For our first application, let us now look at blocks. For any $k \in \mathbb{N}$, a *k-block* in a graph is any maximal set of at least k vertices no two of which can be separated by $< k$ vertices. Every k-block X in a graph G defines an $\mathcal{B}_k$-tangle of S_k for

$$\mathcal{B}_k := \Big\{ \sigma \subseteq \vec{S} : \Big| \bigcap \{ B \mid (A,B) \in \sigma \} \Big| < k \Big\},$$

where S is the set of all separations of G: just orient every separation in S_k towards X. Conversely, given any $\mathcal{B}_k$-tangle τ of S_k in G, the set $\bigcap\{ B \mid (A,B) \in \tau \}$ is a k-block in G, which in turn defines τ as indicated.

One of the earliest applications of what later became $\mathcal{F}$-tangles was that finite graphs have canonical tree-decompositions that distinguish all their k-blocks [7]. This extended a classical decomposition theorem of Tutte for 2-blocks. But there was no tangle-tree duality theorem for blocks, since $\mathcal{B}_k$ does not consist of stars.

It was shown much later that blocks can in fact be captured by $\mathcal{F}$-tangles with $\mathcal{F}$ consisting of stars [12, 13, 17]. But these $\mathcal{F}$ have to be constructed recursively for each given graph, and are therefore highly artificial: tailor-made for the purpose, but not defining $\mathcal{F}$-tangles that can be understood from their definition. As a consequence, the certificates they yield in theory for the non-existence of blocks are of little use in practice. Our Theorem 6.5, however, offers very tangible certificates:

Theorem 7.1. *Let S be the set of separations of a graph G, with the usual order function [11], and let $k \in \mathbb{N}$. Then exactly one of the following assertions holds:*

(i) *G has a k-block;*
(ii) *there exists an $\mathcal{B}_k$-tree of $\vec{S}_k$.*

In the case of (ii), the $\mathcal{B}_k$-tree can be chosen to be irreducible, efficient, and ordered.

Proof. We have already seen that k-blocks induce $\mathcal{B}_k$-tangles of S_k and vice versa. We can thus replace (i) with the assertion that S_k has a $\mathcal{B}_k$-tangle, and proceed to apply Theorem 6.5.

Let us check the theorem's premise. The set $\mathcal{B}_k$ is standard for $\vec{S}_k$, since the all co-trivial oriented separations in $\vec{S}_k$ are of the form (V, A), where $V = V(G)$ and $|A| < k$.[23] Clearly, then, $\{\{(V, A)\}\} \in \mathcal{B}_k$.

For a proof that $\mathcal{B}_k$ is rich for $\vec{S}_k$, it suffices by Lemma 4.4 to show that $\mathcal{B}_k$ is closed under minimization in $\vec{S}$. Given any $\sigma \in \mathcal{F}$, let σ' be obtained from σ by replacing every $(A, B) \in \sigma$ by some $(A', B') \leq (A, B)$ from $\vec{S}$.[24] As

$$\bigcap \{ B' : (A', B') \in \sigma' \} \subseteq \bigcap \{ B : (A, B) \in \sigma \},$$

our assumption of $\sigma \in \mathcal{B}_k$ clearly implies that $\sigma' \in \mathcal{B}_k$, as required. □

Next, profiles. Profiles of graphs were introduced by Hundertmark [20], who noticed that the above-mentioned 'tree-of-tangles' theorem for blocks [7] required in its proof much less information about blocks than is provided in their definition. All that is needed is that the orientation of S_k which a k-block induces[25] avoids the set $\mathcal{F}$ of triples of the form $\{\vec{r}, \vec{s}, \overleftarrow{r} \vee \overleftarrow{s}\}$. In the context of graphs, this supremum is taken in the universe of all the separations of the given graph.

More generally, given any separation system $\vec{S}$ in some universe $\vec{U}$ of separations, we let

$$\mathcal{P} = \{ \{\vec{r}, \vec{s}, \overleftarrow{r} \vee \overleftarrow{s}\} : \vec{r}, \vec{s} \in \vec{U} \},$$

and refer to the $\mathcal{P}$-tangles of S as its *profiles*. A profile is *regular* if none of its elements is small.

Hundertmark's discovery was seminal: since graph tangles are profiles too, the proof of the canonical tree-of-tangles theorem for blocks [7], once rewritten for profiles [14], yielded *canonical* trees of graph tangles – something that Robertson and Seymour [21] had been unable to achieve directly by their tangle methods.

So what about tangle-tree duality for profiles? It was the hope to find such a dichotomy theorem for profiles that motivated the original generalization of graph tangles to $\mathcal{F}$-tangles in abstract separation systems [16], since profiles had this form by definition. But the hope proved elusive. We did indeed find a duality theorem for $\mathcal{F}$-tangles in [16], but this once more required that $\mathcal{F}$ consisted of stars. This was good enough for tangles, since $\mathcal{T}$-tangles coincide with $\mathcal{T}^*$-tangles, as well as for some further duality theorems in graphs and matroids [15], but there was no tangle-tree dichotomy theorem for profiles among these.

As in the case of blocks it is possible to construct, recursively, sets $\mathcal{F}$ of stars in submodular separation systems embedded in a universe so that its profiles are precisely its $\mathcal{F}$-tangles. One can then apply [16] to obtain a tangle-tree duality theorem for profiles, encoded as such $\mathcal{F}$-tangles [12].

Alternatively, we can now obtain certificates for the non-existence of profiles directly from Theorem 6.5.

As preparation, we need a lemma. Given a separation system $\vec{S}$ in some universe $\vec{U}$, we call a profile of S *strong* if it has no subset in

$$\mathcal{P}_s = \{ \{\vec{r}, \vec{s}, \overleftarrow{t}\} : \vec{r}, \vec{s}, \vec{t} \in \vec{U} \text{ and } \overleftarrow{t} \leq \overleftarrow{r} \vee \overleftarrow{s} \}.$$

Note that strong profiles are regular: if $\vec{r} \leq \overleftarrow{r}$ then no strong profile contains $\vec{r}$, since $\vec{r} \leq \overleftarrow{r} \vee \overleftarrow{r}$ puts $\{\vec{r}\}$ in $\mathcal{P}_s$. If $\vec{S}$ is submodular and $\vec{U}$ is distributive, the converse holds too:

[23] In addition, the set A has to lie in the separator of another separation of order $< k$.

[24] Recall that, for graph separations, $(A', B') \leq (A, B)$ if and only if $A' \supseteq A$ and $B' \subseteq B$.

[25] orient every separation towards the side that contains the given k-block

Lemma 7.2. *Every regular profile of a submodular separation system in a distributive universe of separations is strong.*

Proof. Let $(\vec{S}, \leq)$ be any submodular separation system in a distributive lattice $\vec{U}$. Suppose S has a regular profile τ that is not strong. Then there exist $\vec{r}, \vec{s}, \overleftarrow{t}$ in τ such that $\overleftarrow{t} \leqslant \overleftarrow{r} \vee \overleftarrow{s}$. Choose this triple with $\overleftarrow{t}$ as large as possible under $\leq$.

Our aim will be to show that both $\vec{t} \vee \vec{r}$ and $\vec{t} \vee \vec{s}$ lie in $\vec{S}$. For if they do, they must lie in τ by the consistency of τ, since $\vec{t} \vee \vec{r} \geq \vec{r} \in \tau$ and $\vec{t} \vee \vec{s} \geq \vec{s} \in \tau$ and τ is regular.[26] Then, as

$$(\vec{t} \vee \vec{s})^* \vee (\vec{t} \vee \vec{r})^* = (\overleftarrow{t} \wedge \overleftarrow{s}) \vee (\overleftarrow{t} \wedge \overleftarrow{r}) = \overleftarrow{t} \wedge (\overleftarrow{s} \vee \overleftarrow{r}) = \overleftarrow{t}$$

by distributivity, we shall have $\tau \supseteq \{\vec{t} \vee \vec{r}, \vec{t} \vee \vec{s}, \overleftarrow{t}\} \in \mathcal{P}$, which contradicts our assumption that τ is a profile.

So let us show that $\vec{t} \vee \vec{r} \in \vec{S}$; the proof of $\vec{t} \vee \vec{s} \in \vec{S}$ is analogous. By submodularity we have $\vec{t} \vee \vec{r} \in \vec{S}$ as desired unless $\vec{t} \wedge \vec{r} \in \vec{S}$, so let us assume this. Then also $\vec{S} \ni (\vec{t} \wedge \vec{r})^* = \overleftarrow{t} \vee \overleftarrow{r} \geq \overleftarrow{t} \in \tau$, so $\overleftarrow{t} \vee \overleftarrow{r} \in \tau$ by the consistency and regularity[26] of τ.

If $\overleftarrow{t} \vee \overleftarrow{r} = \overleftarrow{t}$, then $\overleftarrow{t} \geq \overleftarrow{r}$ and thus $\vec{t} \leq \vec{r}$, giving $\vec{t} \vee \vec{r} = \vec{r} \in \vec{S}$ as desired. But otherwise $\overleftarrow{t} \vee \overleftarrow{r} > \overleftarrow{t} \in \tau$, which contradicts the maximality of $\overleftarrow{t}$ in our choice of $\vec{r}, \vec{s}, \overleftarrow{t} \in \tau$ with $\overleftarrow{t} \leq \overleftarrow{r} \vee \overleftarrow{s}$, since also $\overleftarrow{t} \vee \overleftarrow{r} \leq \overleftarrow{r} \vee \overleftarrow{s}$ by definition of the supremum $\overleftarrow{t} \vee \overleftarrow{r}$. □

Lemma 7.2 implies that the regular profiles of a submodular separation system in a distibutive universe are precisely its $\mathcal{P}_{\mathrm{s}}$-tangles. Indeed, being profiles, they are $\mathcal{P}$-tangles and therefore consistent, and they avoid $\mathcal{P}_{\mathrm{s}}$ since, by the lemma, they are strong, so they are $\mathcal{P}_{\mathrm{s}}$-tangles. Conversely, every $\mathcal{P}_{\mathrm{s}}$-tangle avoids $\mathcal{P} \subseteq \mathcal{P}_{\mathrm{s}}$ and is therefore a profile, and it is strong since $\mathcal{P}_{\mathrm{s}}$-tangles cannot contain small separations, as noted earlier.

Theorem 7.3. *Let $\vec{S}$ be any submodular separation system in a distibutive universe of separations. Then exactly one of the following assertions holds:*

(i) *there exists a regular profile of S;*
(ii) *there exists an $\mathcal{P}_{\mathrm{s}}$-tree of $\vec{S}$.*

In the case of (ii), the $\mathcal{P}_{\mathrm{s}}$-tree can be chosen to be irreducible, efficient, and ordered with respect to any order function on S.

Proof. Our aim is us apply Theorem 6.5 to $\vec{S}$ with $\mathcal{F} = \mathcal{P}_{\mathrm{s}}$, and richness referring to any order function on S we may choose. As remarked after Lemma 7.2, the regular profiles of S are precisely its $\mathcal{P}_{\mathrm{s}}$-tangles. It remains to check that $\mathcal{P}_{\mathrm{s}}$ is standard and rich for $\vec{S}$.

For a proof that $\mathcal{P}_{\mathrm{s}}$ is standard, we have to show that $\{\overleftarrow{s}\} \in \mathcal{P}_{\mathrm{s}}$ whenever $\vec{s}$ is trivial in $\vec{S}$. Let $r \in S$ witness that $\vec{s}$ is trivial. Then $\vec{r}, \overleftarrow{r} < \vec{s}$, so $\overleftarrow{s} < \vec{r} < \vec{s}$. Hence $\overleftarrow{s} \leq \vec{s} \vee \vec{s}$, which puts $\{\overleftarrow{s}\}$ in $\mathcal{P}_{\mathrm{s}}$.

For a proof that $\mathcal{P}_{\mathrm{s}}$ is rich for $\vec{S}$, it suffices by Lemma 4.4 to show that $\mathcal{P}_{\mathrm{s}}$ is closed under minimization in $\vec{S}$. Given any set $\sigma = \{\vec{r}, \vec{s}, \overleftarrow{t}\} \in \mathcal{P}_{\mathrm{s}}$, consider any set $\sigma' = \{\vec{r'}, \vec{s'}, \overleftarrow{t'}\}$ with $\vec{r'} \leq \vec{r}$ and $\vec{s'} \leq \vec{s}$ and $\overleftarrow{t'} \leq \overleftarrow{t}$. Then $\sigma' \in \mathcal{P}_{\mathrm{s}}$ as desired, since $\overleftarrow{t'} \leq \overleftarrow{t} \leq \overleftarrow{r} \vee \overleftarrow{s} \leq \overleftarrow{r'} \vee \overleftarrow{s'}$; the middle inequality comes from $\sigma \in \mathcal{P}_{\mathrm{s}}$. □

[26] Compare Lemma 2.1 and the remark preceding it.

Finally, let us apply our results to tangles defined by clusters in large datasets. The basic setup is that $\vec{S}$ is the set 2^V of all subsets of our data set V, the involution $*$ given by complementation in V. Then S is the set of bipartitions $s = \{A, B\}$ of V, whose orientations are $\vec{s} = A$ and $\overleftarrow{s} = B$ (say). We then choose an order function on S that assigns high values to separations that divide many pairs of elements of V which we consider as 'close',[27] the idea being that separations of low order cannot divide many close pairs and therefore cannot cut through a cluster, only chip off a few points. This, then, implies that large clusters induce $\mathcal{F}$-tangles for $\mathcal{F}$ defined as

$$\mathcal{C}_n := \big\{ \{ \vec{r}, \vec{s}, \vec{t} \} : |\vec{r} \cap \vec{s} \cap \vec{t}| < n \big\},$$

where the *agreement* value n is chosen to fit the context.

Thus, $\mathcal{C}_n$-tangles in datasets are a bit like $\mathcal{B}_k$-tangles in graphs, except that n is now a value separate from the order k of the tangles considered. An important difference is that, unlike k-blocks in graphs, the clusters captured by those tangles are not normally subsets X of V that we can, or even wish to, specify explicitly: they can be 'fuzzy', and do not have to lie entirely on the side of the separations chosen by the tangle they induce, only mostly. Quantitatively, while the intersection of any three elements of an $\mathcal{C}_n$-tangle contains at least n datapoints, the intersection of all of them will likely be empty – unlike in the case of k-blocks in graphs, which are equal to the intersection of *all* the sides of the graph's separations chosen by the tangle.

In other contexts than generic clustering one sometimes looks for $\mathcal{C}_n$-tangles not of the set S of all the bipartitions of V, but of some specific, hand-designed set S of bipartitions of V. For example, S might be a questionnaire whose elements partition a population V of individuals that have answered it into those that answered yes and those that answered no. In such a context, $\mathcal{C}_n$-tangles can be interpreted as 'typical' ways to answer the questions in S, as *mindsets* found in the population V regarding the questions in S. See [10] for more.

In all these contexts we can now use Theorem 6.1 to efficiently display the clusters or mindsets in V as determined by S, as long as $\mathcal{C}_n$ is standard and rich for $\vec{S}$. If we are interested specifically in certifying that our dataset has no clusters of some desired density at all, we can similarly apply Theorem 6.5.

The sets $\mathcal{C}_n$ are standard for all $\vec{S} \neq \emptyset$ and $n > 0$: then $\{\emptyset\} \in \mathcal{C}_n$, and only $\emptyset$ can be co-trivial in $\vec{S}$. Also, $\mathcal{C}_n$ is clearly closed under minimization in $\vec{S}$: replacing the elements $\vec{s}$ of a triple σ with subsets $\vec{s'} \leq \vec{s}$ leads to $\bigcap \sigma' \subseteq \bigcap \sigma$, so $\sigma \in \mathcal{C}_n$ implies $\sigma' \in \mathcal{C}_n$. By Lemma 4.4, therefore, $\mathcal{C}_n$ is rich for $\vec{S}$.

Theorems 6.1 and 6.5 thus have the following instances in our context:

Corollary 7.4. *Let $\vec{S}$ be any non-empty set of subsets of a set V that is closed under taking complements in V. Let any order function on S be given, and let $n > 0$. Then $\vec{S}$ has an efficient and irreducible ordered $\mathcal{C}_n$-tangle structure tree.* □

Corollary 7.5. *Let $\vec{S}$ be any non-empty set of subsets of a set V that is closed under taking complements in V. Let any order function on S be given, and let $n > 0$. Then exactly one of the following assertions holds:*

(i) *there exists a $\mathcal{C}_n$-tangle of S;*

(ii) *there exists a $\mathcal{C}_n$-tree of $\vec{S}$.*

In case (ii), the $\mathcal{C}_n$-tree can be chosen to be irreducible, efficient, and ordered. □

[27] Choosing this order function to fit the application context allows for considerable variety in using tangles for clustering [10, Chapter 9]. See [3] for a concrete example.

References

[1] Sandra Albrechtsen, *Refining trees of tangles in abstract separation systems: Inessential parts*, Combinatorial Theory **4(1)** (2024; arXiv:2302.01808).

[2] ———, *Optimal trees of tangles: refining the essential parts*, Innovations in Graph Theory **3** (2026), 171–206.

[3] H.v. Bergen and R. Diestel, *Traits and tangles: an analysis of the Big Five paradigm by tangle-based clustering*, J. Math. Psychology **125 (102920)** (2025; arXiv:2411.18670).

[4] ———, *Tangle structure trees II: trees of tangles and tangle-tree duality*, arXiv:2603.17756.

[5] J. Carmesin, R. Diestel, M. Hamann, and F. Hundertmark, *Canonical tree-decompositions of finite graphs I. Existence and algorithms*, J. Comb. Theory Ser. B **116** (2016), 1–24; arXiv:1305.4668.

[6] ———, *Canonical tree-decompositions of finite graphs II. Essential parts*, J. Comb. Theory Ser. B **118** (2016), 268–283; arXiv:1305.4909.

[7] J. Carmesin, R. Diestel, F. Hundertmark, and M. Stein, *Connectivity and tree structure in finite graphs*, Combinatorica **34** (2014), no. 1, 1–35; arXiv:1105.1611.

[8] J. Carmesin and J. Kurkofka, *Entanglements*, J. Comb. Theory, Ser. B **164** (2024), 17–28; arXiv:2205.11488.

[9] R. Diestel, *Abstract separation systems*, Order **35** (2018), 157–170; arXiv:1406.3797.

[10] ———, *Tangles: A structural approach to artificial intelligence in the empirical sciences*, Cambridge University Press, 2024. e-book: `tangles-book.com/book/`.

[11] ———, *Graph theory* (6th edition), Springer-Verlag, 2025.
Electronic edition available at `http://diestel-graph-theory.com/`.

[12] R. Diestel, Ph. Eberenz, and J. Erde, *Duality theorems for blocks and tangles in graphs*, SIAM J. Discrete Math. **31** (2017), no. 3, 1514–1528; arXiv:1605.09139.

[13] R. Diestel, J. Erde, and D. Weißauer, *Structural submodularity and tangles in abstract separation systems*, J. Comb. Theory Ser. A **167C** (2019), 155–180; arXiv:1805.01439.

[14] R. Diestel, F. Hundertmark, and S. Lemanczyk, *Profiles of separations: in graphs, matroids, and beyond*, Combinatorica **39** (2019), no. 1, 37–75; arXiv:1110.6207.

[15] R. Diestel and S. Oum, *Tangle-tree duality in graphs, matroids and beyond*, Combinatorica **39** (2019), 879–910. arXiv:1701.02651.

[16] ———, *Tangle-tree duality in abstract separation systems*, Advances in Mathematics **377** (2021), 107470; arXiv:1701.02509.

[17] Ph. Eberenz, *Characteristics of profiles*, Master's Thesis, Hamburg University, 2015.

[18] C. Elbracht, J. Kneip, and M. Teegen, *Obtaining trees of tangles from tangle-tree duality*, J. Combinatorics **13** (2022), 251–287; arXiv:2011.09758.

[19] Joshua Erde, *Refining a tree-decomposition which distinguishes tangles*, SIAM J. Discrete Math. **31** (2017), 1529–1551; arXiv:1512.02499.

[20] F. Hundertmark, *Profiles. An algebraic approach to combinatorial connectivity*, unpublished manuscript 2011.

[21] N. Robertson and P.D. Seymour, *Graph minors. X. Obstructions to tree-decomposition*, J. Comb. Theory, Series B **52** (1991), 153–190.

TANGLE STRUCTURE TREES, PART II: TREES OF TANGLES AND TANGLE-TREE DUALITY

(HVB & RD) UNSUBMITTED JOURNAL VERSION

Abstract. Tangle structure trees, introduced in [3], offer a unified data structure that displays all the tangles of a graph or data set together with certificates for the non-existence of any other tangles, either locally or overall. In this paper we apply tangle structure trees to derive new versions of the two fundamental tangle theorems: the tree-of-tangles theorem, and the tangle-tree duality theorem.

We extend the tree-of-tangles theorem to $\mathcal{F}$-tangles that need not be profiles. When $\mathcal{F}$ consists of stars of separations, as it does in classical tangle-tree duality theorems, we show how to convert tangle structure trees that certify the non-existence of $\mathcal{F}$-tangles into tree-decompositions that certify this in the way known from graph tangles, as S-trees over $\mathcal{F}$.

1. Introduction

In this sequel to [3] we explore further applications of *tangle structure trees*. These are data structures, introduced in [3], which simultaneously display

- all the tangles of a given graph or abstract separation system; and
- for all its orientations that are not tangles, certificates of why they are not.

Previously, there were two separate types of theorem in tangle theory for these two purposes: *tree-of-tangle* theorems, which display all the tangles of a graph or data set in a tree-like way that 'distinguishes' different tangles by specifying a minimum-order separation that they orient differently; and *tangle-tree duality* theorems that display, also in a tree-like way, certificates for why the various orientations of the separation system are not tangles – typically in the form of small subsets of oriented separations that are *forbidden* in the given type of tangle. These two fundamental theorems, which exist in various forms and degrees of generality, are thought of as the two pillars of both graph and abstract tangle theory.

The unified way in which tangle structure trees serve both these purposes still allows us to extract, if desired, separate corollaries of the two types. In [3] we did this for the second type: we derived new tangle-tree duality-type theorems for tangles whose defining forbidden subsets were not 'stars' of separations, nested sets of typically three oriented separations pointing towards each other. All previously known tangle-tree duality theorems, in particular those for graphs, were for tangles defined by excluding such stars.

Our first main result in this paper is to extract from the tangle structure trees established in [3] a new theorem of the first type: a tree-of-tangles theorem for tangles whose set $\mathcal{F}$ of defining forbidden subsets need not form a 'profile'. All known tree-of-tangles theorems so far are for profiles, in fact, for *robust* profiles.

Date: October 6, 2026.

Robust profiles are still a sweeping generalization of graph tangles. But we shall prove here that we can generalize them further: we shall obtain a tree-of-tangles theorem for all robust $\mathcal{F}$-tangles, no matter what $\mathcal{F}$ is; in particular, these tangles need not be profiles.

Our second main result is about $\mathcal{F}$-tangles whose set $\mathcal{F}$ of forbidden subsets consists of stars of separations, as in all the classical tangle-tree duality theorems. We prove a general conversion theorem, which allows us to extract from our tangle structure tree for $\mathcal{F}$-tangles of a separation system $\vec{S}$, say, an S-tree over $\mathcal{F}$ if S admits no $\mathcal{F}$-tangles. Such an 'S-tree over $\mathcal{F}$' is the tree structure in the most general tangle-tree duality theorems known so far for abstract separation systems (which include graphs), proved in [11]. Our conversion theorem allows us to re-prove the main result of [11] from the premise for the existence of tangle structure trees, which differs from the premise of the tangle-tree duality theorem in [11].

However we show that our premise for the existence of tangle structure trees follows from an only slightly stronger version of the premise in [11], the version used in practice when the duality theorem from [11] is applied [10]. As a consequence, we shall be able in [2] to use our conversion theorem, coupled with the existence theorems for tangle structure trees in [3], to derive all the known applications of the tangle-tree duality theorem in [11] to concrete structures, such as graphs, matroids, or data sets [10].

The paper is organized as follows. In Section 2 we give a summary of tangle basics, just what is technically needed to read this paper. More on graph tangles can be found in [7], more on abstract tangles and their applications in [4–6]. In Section 3 we introduce tangle structure trees and summarize the main results from [3]; see there more motivational background for tangle structure trees.

Section 4 contains our conversion theorem that extracts from a tangle structure tree of a separation system $\vec{S}$ that has no $\mathcal{F}$-tangles an S-tree over $\mathcal{F}$. Section 5 is an interlude on submodular order functions for separations, which is needed in the rest of the paper; we show how a non-injective submodular order function can be extended to an injective one without losing submodularity. The tree-of-tangles theorem for arbitrary robust $\mathcal{F}$-tangles is proved in Section 6. In Section 7 we show how to recover the essence of the tangle-tree duality theorem of [11] from our conversion theorem and our existence theorem for tangle structure trees.

2. Separation systems and their tangles

In this section we introduce the terminology that abstract tangle theory requires, following [4].[1] Readers familiar with [3] may skip this section.

Tangles of graphs are ways of orienting their separations, each towards one of its two sides. Abstract tangles are designed to work in scenarios where there need not be anything to 'separate'. In order to retain our intuition from graphs, however, we continue to refer to the things of which our abstract tangles pick one of two variants (which they will indeed do) as 'separations'. These are defined by noting some key properties of graph separations and making them into axioms, as follows.

A *separation system* $(\vec{S}, \leqslant, ^*)$ is a set $\vec{S}$, whose elements we call *oriented separations*, that comes with a partial ordering $\leqslant$ on $\vec{S}$ and an order-reversing involution

[1] There is one difference to [4]: for historical reasons, the partial ordering on $\vec{S}$ used there is the inverse of ours. So terms like 'large' and 'small', infima and suprema etc, are reversed.

$^*\colon \vec{S} \to \vec{S}$. Thus, for any two elements[2] $\vec{r}, \vec{s}$ of $\vec{S}$ with $\vec{r} \leqslant \vec{s}$ we have $\vec{r}^* \geqslant \vec{s}^*$. We write $\vec{s}^* =: \overleftarrow{s}$, and call $\overleftarrow{s}$ the *inverse* of $\vec{s}$. While we allow formally that $\vec{s} = \overleftarrow{s}$, in which case we call $\vec{s}$ and s *degenerate*, this does not happen often in practice.[3]

If a separation system $\vec{U}$ happens to be a lattice, that is, if there is a supremum $\vec{r} \vee \vec{s}$ and an infimum $\vec{r} \wedge \vec{s}$ in $\vec{U}$ for every two elements $\vec{r}, \vec{s} \in \vec{U}$, we call $\vec{U}$ a *universe* of separations. It is *distributive* if it is distributive as a lattice. A separation system $\vec{S} \subseteq \vec{U}$ is *submodular* if for every two elements of $\vec{S}$ either their infimum or their supremum in $\vec{U}$ also lies in $\vec{S}$.

Very rarely we may have separations $\vec{s} \leq \overleftarrow{s}$; then $\vec{s}$ is *small* and $\overleftarrow{s}$ is *large*.[4] Separation systems without small elements are *regular*. We say that $\vec{s}$ is *trivial* (and $\overleftarrow{s}$ is *co-trivial*) in $\vec{S}$ if there exists a pair of inverse separations $\vec{r}, \overleftarrow{r} < \vec{s}$ in $\vec{S}$. Trivial separations are clearly large, so co-trivial ones are small, but the converse need not hold. See [4] for more on these technicalities if desired.

The set of *unoriented separations* in $(\vec{S}, \leq, ^*)$ is

$$S := \{\{\vec{s}, \overleftarrow{s}\} : \vec{s} \in \vec{S}\}.$$

We call the elements $\vec{s}, \overleftarrow{s}$ of s its *orientations*. An *orientation of* S is a set $\tau \subseteq \vec{S}$ that contains exactly one orientation of every $s \in S$. An orientation of a subset of S is a *partial orientation* of S. We write $\tau(s)$ for the unique orientation of s contained in τ, and say that s *distinguishes* two partial orientations τ, τ' of S if both are defined on s and $\tau(s) \neq \tau'(s)$.

If $\vec{r} \geqslant \vec{s}$ we say that $\vec{r}$ *points towards* s (and that $\overleftarrow{r}$ *points away from* s). We say that $\vec{r}$ *points towards* an oriented separation $\vec{s}$ whenever it points towards s, i.e., if $\vec{r} \geqslant \vec{s}$ or $\vec{r} \geqslant \overleftarrow{s}$, and similarly for 'points away from'. A *star* is a set σ of non-degenerate oriented separations that point towards each other. As is easy to check, this happens if and only if $\vec{r} \geqslant \overleftarrow{s}$ (and hence $\vec{s} \geqslant \overleftarrow{r}$) for all distinct $\vec{r}, \vec{s} \in \sigma$.

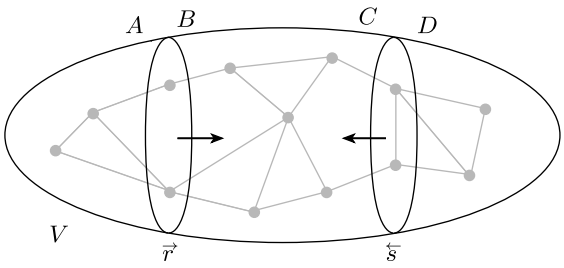


FIGURE 1. Nested separations $r = \{A, B\}$ and $s = \{C, D\}$ of a graph. Their orientations $\vec{r} = (A, B)$ and $\overleftarrow{s} = (D, C)$ point towards each other, since $\vec{r} \geq \vec{s}$ (as $B \supseteq D$) and $\overleftarrow{s} \geq \overleftarrow{r}$ (as $C \supseteq A$).

Two separations $r, s \in S$ are *nested* if they have orientations that are comparable under $\leqslant$; otherwise they *cross*. Oriented separations are *nested* if their underlying unoriented separations are nested. A subset of $\vec{S}$ is *nested* if its elements are pairwise nested.

[2]We often denote the elements of $\vec{S}$ by letters with an arrow, in either direction, precisely in order to have a simple way to refer to their dual elements: by reversing the arrow. But the arrow directions have no meaning: an arbitrary element of $\vec{S}$ could be denoted equally as $\vec{s}$ or as $\overleftarrow{s}$.

[3]The only degenerate separation of a graph $G = (V, E)$, for example, is $\{V, V\}$.

[4]The small separations of a graph G are those of the form (V, A) with $A \subseteq V = V(G)$.

A subset of $\vec{S}$ is *consistent* if no pair of its elements $\vec{r}, \vec{s}$ with $r \neq s$ point away from each other. Stars are examples of consistent nested sets of oriented separations.

We shall often be interested in consistent orientations of S. For each of its elements $\vec{r}$, a consistent orientation of S will also contain every $\vec{s} > \vec{r}$ other than, possibly, $\vec{s} = \overleftarrow{r}$. Consistent partial orientations of S are easily seen to extend to consistent orientations of S, unless they contain a separation that is co-trivial in $\vec{S}$; see [4].

If $\sigma \subseteq \vec{S}$ is consistent, we say that $\vec{s} \in \vec{S}$ is *required* by σ if $\vec{s} \notin \sigma$ and $\sigma \cup \{\overleftarrow{s}\}$ is inconsistent. We shall see in Lemma 2.1 that, pathological cases aside, $\sigma \cup \{\vec{s}\}$ will then be consistent. The *closure* of σ is

$$\lfloor \sigma \rfloor := \sigma \cup \{\vec{s} \in \vec{S} : \vec{s} \text{ is required by } \sigma\}.$$

Note that σ requires $\vec{s} \notin \sigma$ if and only if there exists an $\vec{r} \in \sigma$ such that $r \neq s$ and $\vec{r} < \vec{s}$. Thus

$$\lfloor \sigma \rfloor = \sigma \cup \{ \vec{s} \in \vec{S} : \exists\, \vec{r} \in \sigma \text{ such that } r \neq s \text{ and } \vec{r} < \vec{s}\},$$

which motivates the term of (upward) 'closure'. If $\lfloor \sigma \rfloor$ is consistent, then $\lfloor \lfloor \sigma \rfloor \rfloor$ is (defined and) easily shown to be equal to $\lfloor \sigma \rfloor$.

If σ contains no small separations, the expression above simplifies to

$$\lfloor \sigma \rfloor = \{ \vec{s} \in \vec{S} : \exists\, \vec{r} \in \sigma \text{ such that } \vec{r} \leq \vec{s}\}.$$

Indeed, any $\vec{s}$ in this latter set either lies in σ or there is some $\vec{r} \in \sigma$ such that $\vec{r} < \vec{s}$; in that case $r \neq s$, since otherwise $\overleftarrow{s} = \vec{r} < \vec{s}$, making $\vec{r} \in \sigma$ small.

Lemma 2.1 ([3])**.** *Let $\sigma \subseteq \tau \subseteq \vec{S}$ be consistent sets. Then $\lfloor \sigma \rfloor \subseteq \lfloor \tau \rfloor$. If τ is an orientation of all of S, then $\lfloor \tau \rfloor = \tau$. If τ has no elements that are co-trivial in $\vec{S}$, then $\lfloor \tau \rfloor$ is consistent and $\lfloor \tau \rfloor \setminus \tau$ contains at most one orientation of any $s \in S$.*

Consistent orientations of S can contain small separations. But examples are rare and can be counter-intuitive, so we often exclude them. Co-trivial separations cannot occur in consistent orientations of S. Indeed if $\overleftarrow{s}$ is co-trivial, witnessed by $r \in S$, then every orientation of S will have to orient r too. But it cannot do so consistently with $\overleftarrow{s}$, since both $\vec{r}$ and $\overleftarrow{r}$ are inconsistent with $\overleftarrow{s}$. Similarly, $\lfloor \{\overleftarrow{s}\} \rfloor$ is then inconsistent, since it contains $\vec{r}$ and $\overleftarrow{r}$, which are both inconsistent with $\overleftarrow{s}$.

An *order function* on S is any map $S \to \mathbb{R}$. Unless otherwise mentioned, we denote such order functions as $s \mapsto |s|$. We extend them to $\vec{S}$ by letting $|\vec{s}| := |\overleftarrow{s}| := |s|$. A separation system $(\vec{S}, \leq, ^*)$ given with an order function on S is an *ordered separation system*. For every $k \in \mathbb{R}$ we let

$$\vec{S}_k := \{\vec{s} \in \vec{S} : |s| < k\};$$

this is again an ordered separation system.

We are sometimes interested in orientations of S that do not have certain subsets. We typically collect those together in some set $\mathcal{F}$, whose elements we call *forbidden subsets*. Formally, if $\mathcal{F}$ is any set, we say that $\tau \subseteq \vec{S}$ *avoids* $\mathcal{F}$ if τ has no subset in $\mathcal{F}$, i.e., if no subset of τ is an element of $\mathcal{F}$.

Definition 2.2. An *$\mathcal{F}$-tangle* of S is an $\mathcal{F}$-avoiding consistent orientation of S. The $\mathcal{F}$-tangles of the subsets S_k of S are the $\mathcal{F}$-tangles *in* $\vec{S}$.

When our choice of $\mathcal{F}$ is clear from the context, or arbitrary, we shall also refer to $\mathcal{F}$-tangles simply as *tangles*.

The tangles of the sets $S_k \subseteq S$ are the *k-tangles of S*, or the *tangles of order k* in $\vec{S}$. When we speak of *maximal* tangles in $\vec{S}$, we refer to their partial ordering as subsets of $\vec{S}$. A maximal $\mathcal{F}$-tangle in $\vec{S}$, thus, is an $\mathcal{F}$-tangle τ of some S_k for which there is no $\mathcal{F}$-tangle $\tau' \neq \tau$ of any S_ℓ such that $\tau = \tau' \cap \vec{S}_k$.

3. Tangle structure trees

We adopt the graph-theoretic terminology of [7]. Given a tree T with a root r, we write $\leq_r$ for the associated partial ordering on the nodes of T. Maximal elements, including the root if $|T| = 1$, are called *leaves*. Any direct successors in $\leq_r$ of a node of T are its *children*. We write E_v for the set of edges of T from a node v to its children.

Let $(\vec{S}, \leq, ^*)$ be a separation system, and let $\mathcal{F}$ be any set. A *separation tree* (T, r, β) on $\vec{S}$ consists of a rooted tree (T, r) and an edge labelling $\beta\colon E(T) \to \vec{S}$ such that for every non-leaf $v \in V(T)$ there exists a separation $s_v \in S$ such that β restricts to a bijection $E_v \to \{\vec{s_v}, \overleftarrow{s_v}\}$ and $s_u \neq s_v$ whenever $u <_r v$.

Thus, every non-leaf node v of a separation tree has either one or two children. If there is no need to refer to r or β explicitly, we usually abbreviate (T, r, β) to T. For every node v, we write β_v for the set $\beta\big(E(rTv)\big)$ of edge labels on the path rTv. A separation tree is *consistent* if $\beta_v \subseteq \vec{S}$ is consistent for every node $v \in T$.

A leaf ℓ of a consistent separation tree on $\vec{S}$ is an *$\mathcal{F}$-tangle leaf* if the closure $\lfloor \beta_\ell \rfloor$ of β_ℓ is an $\mathcal{F}$-tangle of S. A non-leaf node v is an *$\mathcal{F}$-tangle node* if there is an $\mathcal{F}$-tangle leaf $\ell \geq_r w$ for every (one or two) successor w of v in $<_r$. Tangle nodes with two children, thus, are precisely the $<_r$-infima of pairs of distinct tangle leaves.

A leaf ℓ is *forbidden* (by $\mathcal{F}$) if β_ℓ contains an element of $\mathcal{F}$ as a subset. Note that tangle leaves are never forbidden. An *$\mathcal{F}$-tangle structure tree of $\vec{S}$* is a consistent separation tree on $\vec{S}$ in which every leaf is either a tangle leaf or forbidden, and for every non-leaf node v the set β_v has no subset in $\mathcal{F}$.

Theorem 3.1 ([3, Theorem 3.7])**.** *Let $\vec{S}$ be a separation system, let $\mathcal{F}$ be any set, and let T be an $\mathcal{F}$-tangle structure tree of $\vec{S}$. Then for every $\mathcal{F}$-tangle τ of S there is a unique leaf $\ell =: \ell(\tau)$ of T with $\lfloor \beta_\ell \rfloor = \tau$. In particular, if all the leaves of T are forbidden, then S has no $\mathcal{F}$-tangle.*

$\mathcal{F}$-tangle structure trees will be our main tool, so let us see when they exist.

Essentially, two conditions on $\mathcal{F}$ are needed to ensure that $\vec{S}$ has an $\mathcal{F}$-tangle structure tree. The first is that $\{\overleftarrow{s}\} \in \mathcal{F}$ for every $\vec{s} \in \vec{S}$ that is trivial in $\vec{S}$. This holds for all relevant $\mathcal{F}$ (see [3] for why), and if it does we call $\mathcal{F}$ *standard for* $\vec{S}$.

The second condition is the one that bites. If T is any separation tree on $\vec{S}$, then every orientation τ of S contains a set β_ℓ for a unique leaf ℓ of T [3, Lemma 3.2]. If τ is consistent and $\lfloor \beta_\ell \rfloor$ orients all of S, then $\lfloor \beta_\ell \rfloor = \tau$ by Lemma 2.1. Hence if τ is not an $\mathcal{F}$-tangle, then $\lfloor \beta_\ell \rfloor$ has a subset σ in $\mathcal{F}$ that witnesses this.

If T is even an $\mathcal{F}$-tangle structure tree, then ℓ is a forbidden leaf, and so we can find such a subset $\sigma \in \mathcal{F}$ of τ not just in $\lfloor \beta_\ell \rfloor$ but even in β_ℓ: among the edge labels of T. In order for this to be possible, we therefore need to make some 'richness' assumption about $\mathcal{F}$: an assumption which, in our example, ensures that $\mathcal{F}$ has enough elements to contain a subset also of β_ℓ as soon as it contains a subset of $\lfloor \beta_\ell \rfloor$.

The main contribution of [3] was to identify such a richness condition on $\mathcal{F}$ that is, essentially, both necessary and sufficient for $\mathcal{F}$-tangle structure trees to exist.

To motivate it, let us first state a stronger condition that is less technical, and which clearly implies that $\mathcal{F}$ contains a subset of β_ℓ whenever it contains a subset of $\lfloor\beta_\ell\rfloor$. This is that $\mathcal{F}$ is *closed under minimization* in every consistent orientation τ of S: that it contains every set $\sigma' \subseteq \tau$ obtained from some $\sigma \subseteq \tau$ in $\mathcal{F}$ by replacing every $\vec{s} \in \sigma$ with some $\vec{s'} \leq \vec{s}$ from τ. Our richness notion in Definition 3.2 will be weaker than this ([3, Lemma 4.4]), but still strong enough to ensure the existence of $\mathcal{F}$-tangle structure trees.

A separation tree T on an ordered separation system $\vec{S}$ is *ordered* if $|s_v| \leqslant |s_w|$ whenever v and w are non-leaves of T with $v \leqslant w$. It is *thoroughly ordered* (in $\vec{S}$) if, for every non-leaf node v, the separation s_v is not oriented by $\lfloor\beta_v\rfloor$ (i.e., $\lfloor\beta_v\rfloor$ contains neither $\vec{s_v}$ nor $\overleftarrow{s_v}$) and has minimum order among the separations not oriented by $\lfloor\beta_v\rfloor$. Every thoroughly ordered separation tree is ordered [3, Lemma 4.1].

We say that $\vec{s} \in \vec{S}$ is *weakly eclipsed* by $\vec{r} \in \vec{S}$ if $\vec{r} < \vec{s}$ and $|r| \leqslant |s|$, and *eclipsed* by $\vec{r}$ if $\vec{r} < \vec{s}$ and $|r| < |s|$. Given any set $\tau \subseteq \vec{S}$, a subset $\sigma \subseteq \tau$ is *efficient* (in τ) if no element of σ is eclipsed by any other element of τ. It is *strongly efficient* if no element of σ is weakly eclipsed by any other element of τ. Our separation tree T is *efficient* if for every leaf ℓ the set β_ℓ is efficient in $\lfloor\beta_\ell\rfloor$. Thoroughly ordered separation trees are efficient [3, Lemma 4.9].

Definition 3.2. A set $\mathcal{F}$ is *rich for* $\vec{S}$ if every consistent orientation of S that has a subset in $\mathcal{F}$ also has a strongly efficient[5] subset in $\mathcal{F}$.

We have the following two theorems for the existence of $\mathcal{F}$-tangle structure trees.

Theorem 3.3 ([3, Theorems 4.6, 4.8])**.** *Let $\vec{S}$ be an ordered separation system, and let $\mathcal{F}$ be any set that is rich and standard for $\vec{S}$. Then there exists a thoroughly ordered $\mathcal{F}$-tangle structure tree of $\vec{S}$. It is unique if the order function on S is injective.*

Note that the uniqueness part of Theorem 3.3 does not require the existence part: an injectively ordered separation system clearly has at most one thoroughly ordered $\mathcal{F}$-tangle structure tree, regardless of what $\mathcal{F}$ is. In particular, $\mathcal{F}$ need not be standard or rich. But we know from [3, Theorem 4.8] that $\mathcal{F}$, assuming it is standard, will be rich for $\vec{S}$ if the order function on S is injective and an $\mathcal{F}$-tangle structure tree exists.

Thoroughly ordered structure trees can be large. But there is a way to reduce them. Given an $\mathcal{F}$-tangle structure tree (T, r, β), consider an edge vw of T such that w is a child of v. We call vw *necessary* for a tangle leaf $\ell \geq w$ of T if $\beta(vw)$ is a $\leq$-minimal element of the tangle $\lfloor\beta_\ell\rfloor$, or equivalently, of β_ℓ. We call vw *necessary* for a forbidden leaf $\ell \geq w$ if every subset of β_ℓ in $\mathcal{F}$ contains $\beta(vw)$. A node v of T is *necessary in* T if for every child w of v there exists a leaf $\ell \geqslant w$ such that vw is necessary for ℓ. If every node of T is necessary, then T is *irreducible*.

It is easy to make a tangle structure tree irreducible by contracting edges at unnecessary nodes and deleting certain subtrees [3, Section 5]. This preserves efficiency and being ordered, but not being thoroughly ordered. Our second existence theorem for $\mathcal{F}$-tangle structure trees gains efficiency and irreducibility over that from Theorem 3.3 at the expense of losing the thoroughness of its ordering:

Theorem 3.4 ([3, Theorem 6.1])**.** *Let $\vec{S}$ be an ordered separation system, and let $\mathcal{F}$ be any set that is rich and standard for $\vec{S}$. Then $\vec{S}$ has an efficient and irreducible ordered $\mathcal{F}$-tangle structure tree.*

[5] in this orientation of S

See [3, Theorem 6.2] for a long list of properties of these $\mathcal{F}$-tangle structure trees that make them useful for tangle analysis, both in theory and in applications.

4. Tangle-tree duality

If we apply Theorem 3.4 to a separation system $\vec{S}$ that has no $\mathcal{F}$-tangles, it tells us that there is an $\mathcal{F}$-tangle structure tree all whose leaves are forbidden. Let us call such structure trees *$\mathcal{F}$-trees*. Conversely, if there is an $\mathcal{F}$-tree of $\vec{S}$ then S has no $\mathcal{F}$-tangle, by Theorem 3.1.

Theorem 3.4 thus implies the following dichotomy:

Theorem 4.1 ([3, Theorem 6.4])**.** *Let $\vec{S}$ be an ordered separation system, and let $\mathcal{F}$ be standard and rich for $\vec{S}$. Then exactly one of the following assertions holds:*

(i) *there exists an $\mathcal{F}$-tangle of S;*

(ii) *there exists an $\mathcal{F}$-tree of $\vec{S}$.*

In the case of (ii), the $\mathcal{F}$-tree can be chosen to be ordered, irreducible and efficient.

Note that, unlike in all known theorems of this type, the set $\mathcal{F}$ of forbidden subsets need not consist of stars of separations (see Section 2).

In the special case that $\mathcal{F}$ does consist of stars of separations, we shall see in Corollary 4.5 that – as long as the $\mathcal{F}$-tree in (ii) is irreducible, which the theorem says we may require – the separations in $\beta(E(T)) \subseteq \vec{S}$ will be nested. In contexts where separation systems describe genuine separations, of a set or of a graph, say, nested sets of separations are known to cut up the structure they separate in a tree-like way – for example, by a tree-decomposition of the graph. In this way, $\mathcal{F}$-trees impose a tree structure on such a set or graph, which in turn implies that its separations we are considering have no $\mathcal{F}$-tangle.

This phenomenon is one of the corner stones of classical tangle theory, known as *tangle-tree duality*. Our aim now is to derive such a dichotomy theorem also in our context, as an application of our more comprehensive tangle structure trees. Our version, Corollary 4.6 below, will not imply the classical tangle-tree duality theorem for abstract separation systems [11], because our richness assumptions made for $\mathcal{F}$ are different. But it will offer the same dichotomy. And we shall see in [2] that it implies all known instances of the classical theorem for concrete structures, such as graphs or matroids.

To carry this out formally, let us introduce the terms in which classical tangle-tree duality is cast. Let $\vec{E}(T)$ denote the set of orientations of the edges of a tree T. For an oriented edge $\vec{e} = (t, t')$ of T we call $t =: i(\vec{e})$ its *initial node* and $t' =: t(\vec{e})$ its *terminal node*, and we denote its inverse (t', t) by $\overleftarrow{e}$. We think of $t(\vec{e})$ as the node that $\vec{e}$ *points towards*. There is a natural partial ordering $\leq$ on $\vec{E}(T)$ in which $\vec{e} > \vec{f}$ if $f \neq e$ and the unique path in T from e to f starts at $t(\vec{e})$ and ends at $i(\vec{f})$. With the involution $^*\colon \vec{e} \mapsto \overleftarrow{e}$, this makes $(\vec{E}(T), \leq, ^*)$ into a separation system.

Since every two edges e, f of T have comparable orientations, the elements of $\vec{E}(T)$ are all nested. And for every node t of T, the set

$$\vec{F}_t := \{\, \vec{e} \in \vec{E}(T) : t(\vec{e}) = t \,\}$$

is a star in this separation system.

Now let $(\vec{S}, \leq, ^*)$ be any other separation system – typically, but not necessarily, consisting of separations of a set or graph. An *S-tree* (T, α) is a tree T together

with a map $\alpha\colon \vec{E}(T) \to \vec{S}$ that satisfies $\alpha(\overleftarrow{e}) = \alpha(\vec{e})^*$ for every $\vec{e} \in \vec{E}(T)$. It is an S-tree *over* a set $\mathcal{F}$ if $\alpha(\vec{F}_t) \in \mathcal{F}$ for every $t \in V(T)$.

If α preserves the partial ordering $\leq$ on $\vec{E}(T)$, i.e. if $\vec{e} \leq \vec{f}$ implies $\alpha(\vec{e}) \leq \alpha(\vec{f})$, and the image of α contains no degenerate separations, then the sets $\alpha(\vec{F}_t) \in \mathcal{F}$ will be stars in $\vec{S}$, because the $\vec{F}_t$ are stars in $\vec{E}(T)$. Conversely:

Lemma 4.2 ([5])**.** *Let (T, α) be an S-tree over a set stars. If α is injective on $\vec{F}_t$ for every node t of T, then $\vec{e} \leq \vec{f}$ implies $\alpha(\vec{e}) \leq \alpha(\vec{f})$ for all $\vec{e}, \vec{f} \in \vec{E}(T)$. In particular, $\alpha(\vec{E}(T))$ is a nested subset of $\vec{S}$.*

Conversely, for every nested, regular, finite separation system $(\vec{S}, \leq, ^*)$ there is an S-tree (T, α) over stars such that $\alpha(\vec{E}(T)) = \vec{S}$ and α is injective on the stars $\vec{F}_t$.

We remark that if α is not injective on some $\vec{F}_t$, then $\alpha(\vec{E}(T))$ can fail to be nested.

Lemma 4.2 motivates the study of S-trees over sets $\mathcal{F}$ of stars (of separations): they make precise, in a more general way than tree-decompositions do, the 'tree-like way' in which nested sets of separations cut up any underlying structure.

The existence of an S-tree (T, α) over a set $\mathcal{F}$, not necessarily consisting of stars, precludes the existence of an $\mathcal{F}$-tangle of S. Indeed, given any orientation τ of S, the set $\alpha^{-1}(\tau)$ is an orientation of the edges of T. Since trees contain no cycles, this orientation will contain $\vec{F}_t$ for some t, so that $\alpha(\vec{F}_t) \subseteq \tau$. But as (T, α) is an S-tree over $\mathcal{F}$ we have $\alpha(\vec{F}_t) \in \mathcal{F}$, so τ is not an $\mathcal{F}$-tangle.

The converse implication, that there is an S-tree over $\mathcal{F}$ whenever S has no $\mathcal{F}$-tangle, holds only when $\mathcal{F}$ consists of stars. But assuming this is not enough. In [11], which offers the most general theorem of this type to date, there is the further assumption that $\vec{S}$ must be *$\mathcal{F}$-separable*. This is a richness condition on $\mathcal{F}$, different from ours, which requires $\mathcal{F}$ to also contain stars which its elements 'induce on the sides' of certain separations in S. The precise notion of $\mathcal{F}$-separability is quite technical, but immaterial here, so we refer the reader to [10, 11] for details.

The following dichotomy between $\mathcal{F}$-tangles and S-trees over $\mathcal{F}$ is the classical tangle-tree duality theorem for abstract separation systems:

Theorem 4.3. [11] *Let $\vec{S}$ be a separation system inside a universe $\vec{U}$ of separations, and let $\mathcal{F} \subseteq 2^{\vec{S}}$ be a set of stars that is standard for $\vec{S}$. If $\vec{S}$ is $\mathcal{F}$-separable, then exactly one of the following assertions holds:*

(i) *there exists an $\mathcal{F}$-tangle of S;*
(ii) *there exists an S-tree over $\mathcal{F}$.*

When $\mathcal{F}$ consists of stars, we can compare Theorem 4.3 with what our more general Theorem 4.1 says for such $\mathcal{F}$. Both theorems guarantee that if S admits no $\mathcal{F}$-tangle then it has an easily detectable certificate in the form of a certain tree-like structure: an S-tree over $\mathcal{F}$ in the case of Theorem 4.3, and an $\mathcal{F}$-tree of $\vec{S}$ in the case of our new Theorem 4.1.

The following result shows that our Theorem 4.1 is not only more general, because its $\mathcal{F}$ need not consist of stars, but also stronger than Theorem 4.3, because we can obtain S-trees over $\mathcal{F}$ from $\mathcal{F}$-trees when $\mathcal{F}$ does consist of stars. In fact, *every* irreducible $\mathcal{F}$-tree of $\vec{S}$ then has the property that its edge labels in $\vec{S}$ are precisely the edge labels of some S-tree over $\mathcal{F}$:

Theorem 4.4. *Let $\vec{S}$ be a separation system without trivial elements, $\mathcal{F} \subseteq 2^{\vec{S}}$ a set of stars, and (T, r, β) an irreducible $\mathcal{F}$-tree of $\vec{S}$. Then there exists an S-tree (T', α) over $\mathcal{F}$ and a map $\gamma\colon V(T') \cup \vec{E}(T') \to V(T) \cup E(T)$ such that*

(i) *γ is a bijection from the nodes of T' to the leaves of T;*
(ii) *γ is a bijection from the oriented edges of T' to the edges of T;*
(iii) *for every edge e of T' there is a node v_e of T such that $\{\gamma(\vec{e}), \gamma(\overleftarrow{e})\} = E_{v_e}$, and these v_e are distinct for different e;*
(iv) *for every oriented edge $\vec{e}$ of T', the map γ sends $t(\vec{e})$ to a leaf $\ell > v_e$ of T such that $\gamma(\vec{e})$ is the first edge of the path $v_eT\ell$;*
(v) *$\alpha = \beta \circ \gamma$.*

In particular, $\beta(E(T)) = \alpha(\vec{E}(T'))$.

Before we prove Theorem 4.4, let us note a corollary for $\mathcal{F}$-trees that seems remarkable in its own right, quite independently of tangle-tree duality. If $\mathcal{F}$ consists of stars, the edges of any irreducible $\mathcal{F}$-tree map to nested separations in $\vec{S}$:

Corollary 4.5. *Let $\vec{S}$ be a separation system without trivial separations, and let $\mathcal{F} \subseteq 2^{\vec{S}}$ be a set of stars. Then for every irreducible $\mathcal{F}$-tree (T, r, β) of $\vec{S}$ the image $\beta(E(T))$ of its edges is a nested set of separations in $\vec{S}$.*

Proof. Our aim is to apply Lemma 4.2 to the S-tree (T', α) provided by Theorem 4.4, to show that $\alpha(\vec{E}(T')) = \beta(E(T))$ is nested. To apply the lemma, we have to check that α is injective on the stars $\vec{F}_t$ in $\vec{E}(T')$.

Given a node t of T', we have $t = t(\vec{e})$ for all $\vec{e} \in \vec{F}_t$. For $\ell := \gamma(t)$ we therefore have $\gamma(\vec{F}_t) \subseteq E(rT\ell)$ by Theorem 4.4 (iv). In order to show that α is injective on $\vec{F}_t$, it thus suffices by (v) $\alpha = \beta \circ \gamma$ and the injectivity of γ to show that β is injective on $E(rT\ell)$. This, however, follows from the last clause in the definition of separation trees. □

Proof of Theorem 4.4. We start our construction of T' by taking as $V(T')$ any set we can map to the leaves of T by a bijection γ. This ensures (i).

We shall give T' one edge e for every non-leaf node $v =: v_e$ of T. To choose the ends of e in T', consider the children w_1 and w_2 of v in T.[6] For each $i \in \{1, 2\}$, since T is irreducible, the edge vw_i is necessary for some leaf $\ell_i \geq w_i$ of T.[7] Let e join the nodes $\gamma^{-1}(\ell_1)$ and $\gamma^{-1}(\ell_2)$ of T', and let $\gamma(\vec{e}) := vw_i$ for the orientation $\vec{e}$ of e for which $\gamma(t(\vec{e})) = \ell_i$. This satisfies (ii)–(iv), if we think of T' as a multigraph for now, i.e., allow multiple edges.

However, since v_e is the unique infimum of ℓ_1 and ℓ_2 in T, the pair $\gamma^{-1}(\{\ell_1, \ell_2\})$ of nodes in T' is joined only by the edge e: the graph T' has no parallel edges and no loops. Finally, now that γ is fixed, let $\alpha := \beta \circ \gamma$ as in (v).

It remains to show that T' is a tree, and that (T', α) is over $\mathcal{F}$: as $\alpha(\vec{e})^* = \alpha(\overleftarrow{e})$ by definition of γ and α, it will be clear that (T', α) is an S-tree once we know that T' is a tree.

Let us begin by showing that T' is acyclic. Consider two edges e, f of T' with $t(\vec{e}) = i(\vec{f})$. By (iv), the edges $\gamma(\vec{e})$ and $\gamma(\overleftarrow{f})$ of T at v_e and v_f, respectively, both lie on the path $rT\ell$ for $\ell = \gamma(t(\vec{e})) = \gamma(i(\vec{f}))$. As $e \neq f$, the nodes v_e and v_f of $rT\ell$ are distinct, by (iii), so $v_e < v_f$ or $v_f < v_e$. As $s_u \neq s_v$ for $u < v$ in any separation tree, $\alpha(\vec{e}) = \beta(\gamma(\vec{e}))$ and $\alpha(\overleftarrow{f}) = \beta(\gamma(\overleftarrow{f}))$, compare (v), are thus orientations of distinct separations $s_{v_e} \neq s_{v_f}$ in S.

[6]These are distinct: s_v cannot be degenerate, because vw_1 is necessary for a forbidden leaf $\ell \geq w_1$ of T, so $\beta(vw_1)$ lies in a star $\sigma \in \mathcal{F}$. Stars do not contain degenerate separations.

[7]Our proof will imply that these ℓ_i are unique, but for the definition it suffices to pick any.

Since we chose $\gamma^{-1}(\ell)$ as $t(\vec{e})$ and as $t(\overleftarrow{f})$ when we picked the ends of the edges e and f in T', the edges $\gamma(\vec{e})$ and $\gamma(\overleftarrow{f})$ of T are necessary for ℓ: their β-values $\alpha(\vec{e})$ and $\alpha(\overleftarrow{f})$ each lie in some $\sigma \in \mathcal{F}$ that is a subset of β_ℓ, and every subset of β_ℓ in $\mathcal{F}$ contains both. Consider any such $\sigma \in \mathcal{F}$. Since σ is a star, we have $\alpha(\vec{e}) \geqslant \alpha(\overleftarrow{f})^* = \alpha(\vec{f})$. As $s_{v_e} \neq s_{v_f}$, the above inequality is strict. The α-values along any oriented path in T' thus strictly decrease. Hence T' contains no (oriented) cycles.

To complete our proof that T' is a tree, it suffices to show that it has one more node than edges. The first of these numbers equals the number of leaves of T; the second equals the number of non-leaves of T. This latter number is one less than the former, since T is a binary tree.

To show that (T', α) is an S-tree over $\mathcal{F}$, consider any node t of T'; we have to show that $\alpha(\vec{F}_t) \in \mathcal{F}$. As earlier, all the edges $\vec{e} \in \vec{F}_t$ are such that $\gamma(\vec{e}) \in rT\ell$ for $\ell := \gamma(t)$ is necessary for ℓ, so $\alpha(\vec{F}_t) = \beta(\gamma(\vec{F}_t)) \subseteq \sigma \subseteq \beta_\ell$ for some $\sigma \in \mathcal{F}$. It remains to show that $\alpha(\vec{F}_t) \supseteq \sigma$.

Suppose there exists an $\vec{s} \in \sigma \setminus \alpha(\vec{F}_t)$. As $\vec{s} \in \sigma \subseteq \beta_\ell = \beta(rT\ell)$, we can find an edge vw in $rT\ell$ with $v < w$ such that $\beta(vw) = \vec{s}$. Let $\vec{f} := \gamma^{-1}(vw)$. Then $\alpha(\vec{f}) = \vec{s}$, and $\gamma(\vec{f}) \in E(rT\ell)$ but $\gamma(\vec{f}) \notin \gamma(\vec{F}_t) \subseteq E(rT\ell)$. In particular, $v_f \neq v_e$ for every $\vec{e} \in \vec{F}_t$, and hence $f \neq e$, so neither $\vec{f}$ or $\overleftarrow{f}$ lies in $\vec{F}_t$. But note that v_f and all these v_e lie on the path $rT\ell$, so they are comparable in the tree-order on $V(T)$.

Since T' is a tree, either $\vec{f}$ or $\overleftarrow{f}$ points towards t in T'. Suppose first that $\vec{f}$ does. Then $\vec{f} \geq \vec{e}$ in $\vec{E}(T')$ for some $\vec{e} \in \vec{F}_t$. As in our acyclicity proof earlier, this implies that $\alpha(\vec{f}) \geqslant \alpha(\vec{e})$. As both $\vec{s} = \alpha(\vec{f})$ and $\alpha(\vec{e})$ lie in the star σ, we have $\alpha(\vec{f}) \geq \alpha(\vec{e})^*$ too. As v_f and v_e are distinct nodes on $rT\ell$, which implies that $s = s_{v_f} \neq s_{v_e} = \{\alpha(\vec{e}), \alpha(\vec{e})^*\}$ since T is a separation tree, $\vec{s} = \alpha(\vec{f})$ is thus trivial in $\vec{S}$ witnessed by s_{v_e}, a contradiction to our assumptions about $\vec{S}$.

Suppose now that the edge $\overleftarrow{f}$ points towards t in T'. Then $\overleftarrow{f} \geqslant \vec{e} \in \vec{F}_t$ and hence $\alpha(\overleftarrow{f})^* \geq \alpha(\vec{e})$, while $s = s_{v_f} \neq s_{v_e} = \{\alpha(\vec{e}), \alpha(\vec{e})^*\}$, as earlier. This makes $\vec{s} = \alpha(\overleftarrow{f})$ and $\alpha(\vec{e})$ inconsistent elements of β_ℓ, a contradiction to the fact that T is a consistent separation tree. □

Let us discuss briefly what we can say if $\vec{S}$ does have trivial elements. The proof given above stands until we have to prove that the S-tree T' we have constructed is over $\mathcal{F}$. If $\vec{r} \in \vec{S}$ is trivial in $\vec{S}$ witnessed by $s \in S$, this can indeed fail. Indeed, assume that $\vec{s} = \alpha(\vec{e})$ with $\gamma(t(\vec{e})) = \ell$ and $\gamma(t(\overleftarrow{e})) = \ell'$. Let v_f be the predecessor of v_e in T, and assume that $\beta(v_f v_e) = \vec{r}$. Then $v_f v_e$ is the first edge of the path in T from v_f to both ℓ and ℓ', so either of these would be eligible as $\gamma(t(\vec{f}))$, where $\vec{f}$ is the orientation of f with $\alpha(\vec{f}) = \vec{r}$, as long as $v_f v_e$ is necessary for ℓ and ℓ', respectively. This will be the case if and only if $\vec{r}$ lies in the stars $\sigma \in \mathcal{F}$ or $\sigma' \in \mathcal{F}$ that contain $\vec{s}$ and $\overleftarrow{s}$, respectively; let us assume it lies in both. (Recall that $\beta^{-1}(\vec{s})$ and $\beta^{-1}(\overleftarrow{s})$ are necessary for ℓ and ℓ', respectively, so $\vec{s} \in \sigma$ and $\overleftarrow{s} \in \sigma'$ for some and for all stars $\sigma \subseteq \beta_\ell$ and $\sigma' \subseteq \beta_{\ell'}$ in $\mathcal{F}$.) However, when we constructed T', we chose only one of ℓ and ℓ' as $\gamma(t(\vec{f}))$, say ℓ. We then have $\vec{r} \in \sigma' \setminus \alpha(\vec{F}_{t'})$ for $t' := \gamma^{-1}(\ell')$.

However, we can easily mend this: just add a new leaf t to T' adjacent to t', with $\alpha(tt') := \vec{r}$. We shall then need that $\{\overleftarrow{r}\} \in \mathcal{F}$, but this will be the case if we assume that $\mathcal{F}$ is standard, as we have to in order to obtain T in the first place. More generally, we can add to T' such new leaves at all nodes t' whose $\alpha(\vec{F}_{t'})$ fails to contain a trivial separation $\vec{r}$ whose triviality is witnessed by some s with $\vec{s} \in \alpha(\vec{F}_{t'})$. The

modified tree will no longer satisfy the detailed clauses in Theorem 4.4, but it will be an S-tree over $\mathcal{F}$ all whose non-trivial labels $\alpha(\vec{s})$ are obtained as in our proof.

Note that the 'tree-likeness' in S-trees over $\mathcal{F}$ and in $\mathcal{F}$-trees, respectively, refers to very different structures: while S-trees over $\mathcal{F}$ in Theorem 4.3 exhibit tree-likeness in the (object divided by the) separation system $\vec{S}$ itself – for example, they define tree-decompositions with parts specified by $\mathcal{F}$ if S consists of separations of a graph – the $\mathcal{F}$-trees in Theorem 4.1 are 'tree-like' only in the formal sense of a decision-tree for how to orient the elements of S. It is all the more remarkable, therefore, that the structuring trees found in S by Theorem 4.3 can be recovered from the $\mathcal{F}$-trees in Theorem 4.1, as shown in Theorem 4.4. We thus have a new tangle-tree duality theorem: one that offers the same dichotomy as Theorem 4.3, but with a different premise.

To state this in optimal form, we need one more definition. Given any set $\mathcal{F} \subseteq 2^{\vec{S}}$ and any order function on S, let $\mathcal{F}_{\text{eff}}$ denote the set of all sets $\sigma \in \mathcal{F}$ that are efficient in their closure $\lfloor \sigma \rfloor$ in $\vec{S}$ (see Section 3). Note that this definition makes sense independently of any tangle structure trees.

Corollary 4.6. *Let $\vec{S}$ be an ordered separation system without trivial separations. Let $\mathcal{F} \subseteq 2^{\vec{S}}$ be a rich set of stars. Then exactly one of the following assertions holds:*

(i) *there exists an $\mathcal{F}$-tangle of S;*
(ii) *there exists an S-tree over $\mathcal{F}$.*

In the case of (ii), the S-tree can be chosen to be over $\mathcal{F}_{\text{eff}}$.

Proof. As indicated after Lemma 4.2, the assertions (i) and (ii) cannot both hold. Let us assume that (i) fails, and show (ii). By Theorem 4.1, there is an $\mathcal{F}$-tree T of $\vec{S}$, which may be chosen irreducible, efficient and ordered. If T is efficient, it will be an $\mathcal{F}_{\text{eff}}$-tree, since any $\sigma \subseteq \beta_\ell$ from $\mathcal{F}$ is then efficient in $\lfloor \beta_\ell \rfloor$ and therefore also efficient in $\lfloor \sigma \rfloor \subseteq \lfloor \beta_\ell \rfloor$. Now (ii) and the final statement of the corollary follow from Theorem 4.4. □

While Theorem 4.3 and Corollary 4.6 proclaim the same tangle-tree dichotomy, the requirements they make on $\mathcal{F}$ in their premise are different: $\mathcal{F}$-separability of $\vec{S}$ and richness of $\mathcal{F}$, respectively. In order to derive Theorem 4.3 formally from Corollary 4.6 (in the case that $\mathcal{F}$ consists of stars) we would need to show that if $\vec{S}$ is $\mathcal{F}$-separable then $\mathcal{F}$ is rich for $\vec{S}$. The condition of $\mathcal{F}$-separability, however, was tailor-made for the proof of Theorem 4.3. It therefore seems unlikely that we can formally derive from it that $\mathcal{F}$ is rich. However, in Section 7 we shall derive this from a condition that is only slightly stronger than $\mathcal{F}$-separability. This stronger condition, in fact, holds and is used in all the known applications of Theorem 4.3, rather than $\mathcal{F}$-separability itself. This will be explored in [2].

5. Submodular order functions

Submodular order functions are a key ingredient to much of tangle theory. We shall need them in this paper too, both for our tree-of-tangles theorems in Section 6 and to retrieve the essence of the tangle-tree duality theorem from [11] in Section 7.

In both these applications we shall need tangle structure trees that are not just ordered but thoroughly ordered. These exist reliably only when our order function

is injective. If it is not – which often happens in theoretical applications[8] – we have to tweak it slightly to make it so. We shall prove in this section that we can do this while preserving its submodularity.

We say that an order function o' on a set S *refines* another order function o on S if

$$o(r) < o(s) \;\Rightarrow\; o'(r) < o'(s) \quad (\forall r, s \in S).$$

Thus, o' refines o if and only if $o'(r) \leq o'(s)$ implies $o(r) \leq o(s)$. But we do not have the same implication for '$<$', since o' need not be constant on the subsets of S whose elements have a common value under o.

Let $\vec{S}$ be a separation system with an order function o, and let $\mathcal{F}$ be any set. Recall that the $\mathcal{F}$-tangles of the subsets

$$S_k = \{\, s \in S \mid o(s) < k \,\}$$

of S are the $\mathcal{F}$-tangles *in* $\vec{S}$ with respect to o.

Lemma 5.1. *If o' refines o, then the tangles in $\vec{S}$ with respect to o are also tangles in $\vec{S}$ with respect to o'.*

Proof. If o' refines o then, for every $k \in \mathbb{R}$ and with S_k defined with respect to o as above, the range $o'(S_k)$ in $\mathbb{R}$ is an initial segment of $o'(S)$. For every k choose $k' \in \mathbb{R}$ larger than every element of $o'(S_k)$ but smaller than every element of $o'(S) \setminus o'(S_k)$. Then S_k can be rewritten as $\{\, s \in S \mid o'(s) < k' \}$, which puts the $\mathcal{F}$-tangles of S_k in $\mathrm{tangle}(\vec{S}, \mathcal{F}, o')$. □

Let us consider a separation system $\vec{U}$ that is a universe: one whose partial ordering $\leq$ makes it into a lattice. Any two $\vec{r}, \vec{s} \in \vec{U}$ thus have a supremum $\vec{r} \vee \vec{s}$ and an infimum $\vec{r} \wedge \vec{s}$ in $\vec{U}$. Let us call a function $u \colon \vec{U} \to \mathbb{R}$ *submodular* if

$$u(\vec{r} \vee \vec{s}) + u(\vec{r} \wedge \vec{s}) \leqslant u(\vec{r}) + u(\vec{s})$$

for all $\vec{r}, \vec{s} \in \vec{U}$, and *structurally submodular* if

$$u(\vec{r} \vee \vec{s}) \leq u(\vec{r}) \quad \text{or} \quad u(\vec{r} \wedge \vec{s}) \leq u(\vec{s}) \tag{$\dagger$}$$

for all $\vec{r}, \vec{s} \in \vec{U}$. Note that submodular functions on $\vec{U}$ are also structurally submodular, but not conversely. Unlike submodularity, the structural submodularity of a real function on $\vec{U}$ is preserved when we compose it with an order isomorphism on $\mathbb{R}$, but it is not preserved by summing functions.

Let us call an order function $|\,|$ on the set U of unoriented separations *submodular* if the function $\vec{s} \mapsto |s|$ it induces on $\vec{U}$ is submodular, and similarly for 'structurally submodular'. As earlier, submodularity of order functions on U is preserved under summing (but not under composition with order isomorphisms), while structural submodularity of order functions on U is preserved under composition with order isomorphisms (but not under summing).

Note also that, given any $k \in \mathbb{R}$, the separation system $\vec{U_k} = \{\vec{s} \in \vec{U} : |s| < k\}$ is submodular, as defined at the start of Section 2, as soon as $|\,|$ is structurally submodular on $\vec{U}$. Conversely, if $\vec{S} \subseteq \vec{U}$ is submodular, one can define a structurally submodular order function on U so that $\vec{S} = \vec{U_k}$. But there need not exist a submodular order function on U with this property; see [1, Example 5.11].

It is not hard to refine a submodular order function to make it injective. Ideally, we would like to do this while preserving its submodularity. We shall prove a little

[8] Order functions used for real-world tangle applications, such as in clustering, tend to be injective. But the standard order function on graph separations, for example, is not.

less than this, which is still good enough for our use of submodularity: that we can refine it to an order function that is still structurally submodular.

Lemma 5.2. *Let $\vec{U}$ be a separation universe with a submodular function u. Then $u'\colon \vec{s} \mapsto u(\overleftarrow{s})$ and $w\colon s \mapsto u(\vec{s}) + u(\overleftarrow{s})$ are also submodular order functions on U.*

Proof. Since u is submodular, we also have

$$u'(\vec{s} \vee \vec{r}) + u'(\vec{s} \wedge \vec{r}) = u(\overleftarrow{s} \wedge \overleftarrow{r}) + u(\overleftarrow{s} \vee \overleftarrow{r}) \leqslant u(\overleftarrow{s}) + u(\overleftarrow{r}) = u'(\vec{s}) + u'(\vec{r}),$$

so u' is submodular. The sum of two submodular functions is clearly submodular, hence so is also $w = u + u'$. □

Lemma 5.3. *Let $\vec{U}$ be a separation system and $\vec{t} \in \vec{U}$. The function $\mathbf{1}_{\vec{t}}\colon \vec{U} \to \{0,1\}$ given by*

$$\mathbf{1}_{\vec{t}}(\vec{s}) = \begin{cases} 0 & \text{if } \vec{s} \leqslant \vec{t}, \\ 1 & \text{otherwise.} \end{cases}$$

is submodular.

Proof. Let $\vec{r}, \vec{s} \in \vec{U}$. We have to show that

$$\mathbf{1}_{\vec{t}}(\vec{r} \wedge \vec{s}) + \mathbf{1}_{\vec{t}}(\vec{r} \vee \vec{s}) \leqslant \mathbf{1}_{\vec{t}}(\vec{r}) + \mathbf{1}_{\vec{t}}(\vec{s}).$$

If the right-hand side of the inequality is 2, the inequality holds trivially. So let us assume that $\mathbf{1}_{\vec{t}}(\vec{r}) = 0$. Then $\vec{r} \wedge \vec{s} \leqslant \vec{r} \leqslant \vec{t}$, so $\mathbf{1}_{\vec{t}}(\vec{r} \wedge \vec{s}) = 0$. Hence the left-hand side of the inequality is at most 1. The inequality thus holds unless the right-hand side is zero, i.e., unless also $\mathbf{1}_{\vec{t}}(\vec{s}) = 0$. But then also $\vec{s} \leqslant \vec{t}$ and thus $\vec{r} \vee \vec{s} \leqslant \vec{t}$, so $\mathbf{1}_{\vec{t}}(\vec{r} \vee \vec{s}) = 0$, making the left-hand side zero too. □

It may help to think of the sums in our next lemma as n-ary integer expansions.

Lemma 5.4. *Let $\vec{U}$ be a separation universe, $\iota\colon \vec{U} \to \{0, \dots, |\vec{U}| - 1\}$ a bijection, and $n \in \mathbb{N}$. Then*

$$\gamma_n^\iota(\vec{s}) := \sum_{\vec{t} \in \vec{U},\, \vec{t} \not\geqslant \vec{s}} n^{\iota(\vec{t})} = \sum_{\vec{t} \in \vec{U}} n^{\iota(\vec{t})} \mathbf{1}_{\vec{t}}(\vec{s}).$$

In particular γ_n^ι is submodular.

Proof. The first part is immediate, as the additional summands in the second sum are all zero. The submodularity of γ_n^ι then follows from Lemma 5.3 and the fact that sums of submodular functions are again submodular. □

Lemma 5.5. *Let $\vec{U}$ be a separation universe with a submodular order function o. Then there exists an injective submodular order function o' on U that refines o.*

Proof. We shall construct a small injective and submodular perturbation $\delta\colon U \to \mathbb{R}$, so that $o' := o + \delta$ is injective and refines o.

First choose $\varepsilon > 0$ as follows. Let $\Delta := \{|o(s) - o(r)| : s, r \in U,\ o(s) \neq o(r)\}$. If $\Delta \neq \emptyset$ set $\varepsilon := \min \Delta > 0$; otherwise set $\varepsilon := 1$. In either case $\varepsilon > 0$.

Fix an arbitrary bijection $\iota\colon \vec{U} \to \{0, \dots, |\vec{U}| - 1\}$ and set $\gamma(s) := \gamma_3^\iota(\vec{s}) + \gamma_3^\iota(\overleftarrow{s})$. By Lemma 5.2 and Lemma 5.4 the function $\gamma\colon U \to \mathbb{R}$ is submodular. We claim γ is injective. Write $\gamma(s)$ in base 3:

$$\gamma(s) = \sum_{k=0}^{|\vec{U}|-1} a_k(s)\, 3^k, \qquad a_k(s) \in \{0,1,2\}.$$

By construction we have $a_k(s) \leqslant 1$ if and only if the separation $\vec{t} = \iota^{-1}(k)$ points towards s. If $\gamma(s) = \gamma(r)$ then their base-3 expansions coincide, hence every $\vec{t} \in \vec{U}$ points towards s if and only if it points towards r. From this we will now deduce $s = r$, showing that γ is injective.

Let $\vec{s_0}$ be any orientation of s. Suppose that, for some $n \in \mathbb{N}$, we have defined an orientation $\vec{s_n}$ of s. Then since $\vec{t} := \vec{s_n}$ points towards s, it also points towards r. So there exists an orientation $\vec{r_n}$ of r with $\vec{s_n} \geqslant \vec{r_n}$. Analogously, given any orientation $\vec{r_n}$ of r we can find an orientation $\vec{s_{n+1}}$ of s with $\vec{r_n} \geqslant \vec{s_{n+1}}$. For any orientation $\vec{s_0}$ of s we can thus construct an infinite chain $\vec{s_0} \geqslant \vec{r_0} \geqslant \vec{s_1} \geqslant \dots$ of, alternately, orientations of s and r. As r and s have at most two orientations each, at most two of these inequalities can be strict. Hence for m large enough we have $\vec{s_m} = \vec{r_m}$, and therefore $s = r$ as desired. This completes our proof that γ is injective.

Note that $0 \leq \gamma(s) < 3^{|\vec{U}|}$ for every s. Define

$$\delta(s) := \left(\tfrac{1}{2}\varepsilon/3^{|\vec{U}|}\right)\gamma(s).$$

Then $0 \leqslant \delta(s) < \varepsilon/2$ for all s. Note that δ is injective and submodular, because γ is and $\frac{1}{2}\varepsilon/3^{|\vec{U}|} > 0$.

Set $o' := o + \delta$. As a sum of submodular functions, o' is submodular. To see that o' is injective, suppose $o'(r) = o'(s)$. Then

$$|o(r) - o(s)| = |\delta(r) - \delta(s)| \leqslant |\delta(r)| + |\delta(s)| < \varepsilon.$$

By the choice of ε we must have $o(s) = o(r)$, and hence $\delta(s) = \delta(r)$. Since δ is injective this gives $s = r$. Finally, o' refines o, because whenever $o(r) > o(s)$ we have

$$o'(r) - o'(s) = o(r) - o(s) + \delta(r) - \delta(s) \geqslant \varepsilon - |\delta(r)| - |\delta(s)| > 0\,,$$

so $o'(r) > o'(s)$. □

An *enumeration* of a set S is any bijection between S and $\{1, \dots, |S|\}$.

Lemma 5.6. *Let $\vec{U}$ be a separation universe with a submodular order function o. Then there exists a structurally submodular enumeration of U that refines o.*

Proof. Compose the injective submodular refinement o' of o provided by Lemma 5.5 with an order isomorphism φ from its range to $\{1, \dots, |U|\}$. Since o', being submodular, is also structurally submodular, and structural submodularity is preserved under composition with order isomorphisms, the enumeration $\varphi \circ o'$ of U is structurally submodular, and like o' it refines o. □

6. Trees of tangles

In Section 4 we derived one of the two fundamental types of theorem about tangles from our tangle structure trees: a so-called tangle-tree duality theorem. In this section we shall do the same for the second pillar of tangle theory: its so-called tree-of-tangles theorems.

Let $\vec{S}$ be an ordered separation system. Recall from Section 2 that a separation *distinguishes* two tangles in $\vec{S}$ if they orient it differently. Let us say that it distinguishes them *optimally*[9] if it distinguishes them but no separation in S of lower

[9]The usual terminology is to say 'efficiently' rather than 'optimally'. We use the latter only in this paper, to avoid confusion with our use of the term 'efficient' for subsets of $\vec{S}$.

order does. A nested set $N \subseteq S$ is a *tree*[10] *of the tangles of* S if every two $\mathcal{F}$-tangles of S (for some given $\mathcal{F}$) are distinguished optimally by a separation in N. Likewise, N is a *tree of the tangles in* $\vec{S}$ if it is nested and every two maximal $\mathcal{F}$-tangles in $\vec{S}$ are distinguished optimally by a separation in N.

Definition 6.1. Let $\vec{S}$ be a separation system contained in a universe $\vec{U}$. An orientation τ of S is *robust* in $\vec{U}$ if for no $s \in U$ there is a triple $\{\vec{r}, \overleftarrow{r} \vee \vec{s}, \overleftarrow{r} \vee \overleftarrow{s}\}$ in τ with $|\overleftarrow{r} \vee \vec{s}|, |\overleftarrow{r} \vee \overleftarrow{s}| < |r|$.[11]

These triples will be illustrated in Figure 2. Note that while $\vec{r}$ can lie in such a triple in τ only if it lies in $\vec{S}$, the separation s need not lie in S: it just gives rise to that triple in $\vec{S}$ together with $\vec{r} \in \vec{S}$.

We shall find in every separation system $\vec{S}$ that lives in some universe with an injective submodular order function a tree of all the robust $\mathcal{F}$-tangles of S, as well as a more comprehensive tree of all the robust $\mathcal{F}$-tangles in $\vec{S}$. Here, $\mathcal{F}$ has to be standard and rich, as usual, but need not satisfy any further requirements. But to ease our terminology, we shall encode the robustness assumption in $\mathcal{F}$, by assuming that $\mathcal{F}$ contains all the triples from Definition 6.1. We can then simply speak of (trees of) $\mathcal{F}$-tangles rather than of robust $\mathcal{F}$-tangles.

The most comprehensive tree-of-tangles theorems known so far are for $\mathcal{F}$-tangles that are not only robust but also *profiles*: consistent orientations of a separation system $\vec{S}$ that contain no triple of the form $\{\vec{r}, \vec{s}, \overleftarrow{r} \vee \overleftarrow{s}\}$, where the supremum is once more taken in some fixed universe of separations containing $\vec{S}$ [9].

Tangles of graphs are robust profiles [7].[12] So the tree-of-tangles theorems for robust profiles of abstract separation systems from [9] generalize the tree-of-tangles theorem for graphs of Robertson and Seymour [12], and our results below will further generalize those from [9] to $\mathcal{F}$-tangles that are robust but need not be profiles.

There is another advance in [9] over [12]: the trees of tangles found in [9] are *canonical*. This means, roughly, that the automorphisms of the given graph or separation system $\vec{S}$ leave the tree of tangles $N \subseteq S$ invariant. This is usually proved by observing that the construction of N depends only on invariants of the graph or separation system rather than, say, of an enumeration of $\vec{S}$ assumed for the sake of the construction, as was the case for the first trees of tangles for graphs in [12]. As a consequence, such a canonical tree of tangles can then be computed by an algorithm that yields the same result regardless of the order in which the elements of the graph or separation system are presented to it.

Our tree-of-tangles theorems for all robust $\mathcal{F}$-tangles, Theorems 6.2 and 6.10 below, requires that the order function given on S is injective. This assumption makes canonicity, as defined above, trivial: unless $|S| = 1$, the identity will be the only automorphism of $\vec{S}$ as an *ordered* separation system. As we saw in Section 5, we can make a given submodular order function injective without losing its submodularity. But this conversion is not canonical: it depends on choices not determined by the

[10]Recall that nested sets of separations of structures such as sets or graphs divide them in a 'tree-like' way [5]; hence the 'tree' in the name.

[11]If this definition appears technical, think of orientations of bipartitions of a set as designating the side they point to as 'big'. Robust orientations of set partitions are then like a finite analogue to ultrafilters: big subsets never split into two small subsets.

[12]Indeed, all tangles of abstract separation systems, as defined in [8], are profiles (see there). If the ambient universe of their separation system is distributive, they are easily seen to be robust. Graph tangles are examples of such abstract tangles in a distributive universe.

invariants of the separation system considered.[13] If we ignore the order function on our separation system and refer in our notion of canonicity only to the unordered system, as a poset with an order-reversion involution, then the tree of tangles we shall obtain in this section will not be canonical: it will depend on the choice of (injective) order function on S, because the (thoroughly ordered) tangle structure tree from which we construct it depends on it.

The following definition allows us to encode the robustness of our tangles in $\mathcal{F}$. Let $\vec{S}$ be a separation system in a universe $\vec{U}$ of separations. Let

$$\mathcal{R}(\vec{U}) := \big\{\{\vec{r}, \overleftarrow{r} \vee \vec{s}, \overleftarrow{r} \vee \overleftarrow{s}\} \subseteq \vec{U} : r, s \in U \text{ and } |\overleftarrow{r} \vee \vec{s}|, |\overleftarrow{r} \vee \overleftarrow{s}| < |r|\big\}.$$

These triples are easily identified in Figure 2, for example.

Here is our basic tree-of-tangles theorem for robust $\mathcal{F}$-tangles with arbitrary $\mathcal{F}$, or equivalently, for all $\mathcal{F}$-tangles with arbitrary $\mathcal{F}$ that include $\mathcal{R}(\vec{U})$:

Theorem 6.2. *Let $\vec{U}$ be a universe of separations with an injective and structurally submodular order function. Let $k \in \mathbb{R}$ and $\vec{S} := \vec{U}_k$. Let $\mathcal{F}$ be any set that is rich and standard for $\vec{S}$ and includes $\mathcal{R}(\vec{U})$. Let (T, r, β) be the unique thoroughly ordered $\mathcal{F}$-tangle structure tree for $\vec{S}$ (cf. Theorem 3.3). Let N be the set of all separations $s_v \in S$ with v a tangle node of T.*

Then N is a tree of the tangles of S; in particular, it is nested. It consists of all the separations in S that distinguish two $\mathcal{F}$-tangles of S optimally.

Proof. We shall build up our proof of Theorem 6.2 as a sequence of a few lemmas. Let $\vec{U}$, k, $\vec{S}$, $\mathcal{F}$, (T, r, β) and N be as stated in the theorem, with an injective and structurally submodular order function on U.

The *corners* of two separations $r, s \in S$ are the four separations in U that have orientations of the form $\vec{r} \wedge \vec{s}$. Figure 2 shows a corner of two bipartitions of a set, where $\vec{r} \wedge \vec{s}$ is the intersection of the subsets to which $\vec{r}$ and $\vec{s}$ point.

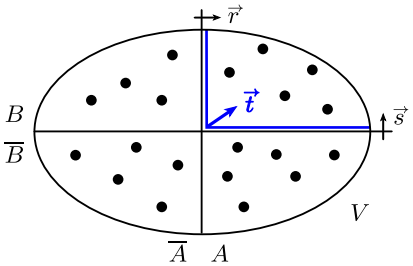


FIGURE 2. A corner $\vec{t} = \vec{r} \wedge \vec{s}$ of two bipartitions of a set V.

Two corners t and t' of r and s lie *on the same side of* r if these separations have orientations $\vec{r} \geq \vec{t}, \vec{t'}$. Corners on the same side of one of two crossing separations are *adjacent*; non-adjacent corners are *opposite*. Since our order function on U is structurally submodular, all four ordered pairs (t, t') of opposite corners satisfy $|t| \leq |r|$ or $|t'| \leq |s|$.

In Figure 2, the triples $\{\vec{r}, \overleftarrow{r} \vee \vec{s}, \overleftarrow{r} \vee \overleftarrow{s}\}$ from the definition of $\mathcal{R}(\vec{U})$ consist of $\vec{r}$ and the two corners on the right-hand side of r both oriented towards r, the last element of the triple being $\overleftarrow{t} = \overleftarrow{r} \vee \overleftarrow{s}$.

[13] In our case, this was the choice of the bijection ι in the proof of Lemma 5.5.

Lemma 6.3. *Any two crossing separations $r, s \in S$ have two adjacent corners that either lie on the same side of r and have order $< |r|$ or lie on the same side of s and have order $< |s|$.*

Proof. Let $|r| \leq |s|$. If all four corners have order $< |s|$ we are done; let t_1 be a corner of order $\geq |s|$. As r and s cross we have $r \neq s \neq t_1$, so in fact $|r| < |s| < |t_1|$.

By structural submodularity, the corner t_2 opposite t_1 has order $\leq |r|$, and hence order $< |r|$, since $t_2 \neq r$ as earlier. If the corner t_3 on the same side of r as t_2 also has order $< |r|$ we are done, so we assume not. Then $|t_3| > |r|$, and its opposite corner t_4 satisfies $|t_4| < |s|$. Since t_4 lies on the same side of s as t_2, and $|t_2| < |r| < |s|$, this completes the proof. □

Lemma 6.4. (i) *If $s \in S$ distinguishes two tangles τ, τ' of S optimally then $s = s_v$, where v is the tangle node of T that is the infimum in T of the tangle leaves $\ell = \ell(\tau)$ and $\ell' = \ell(\tau')$.*[14]

(ii) *Conversely, if the tangle node v of T is the infimum in T of two tangle leaves $\ell = \ell(\tau)$ and $\ell' = \ell(\tau')$, then s_v distinguishes τ and τ' optimally.*

Proof. (i) By definition of v, the tangles $\tau = \lfloor \beta_\ell \rfloor$ and $\tau' = \lfloor \beta_{\ell'} \rfloor$ both include $\lfloor \beta_v \rfloor$, but they differ on s_v. Since s distinguishes τ and τ' optimally, we thus have $|s| \leq |s_v|$, while s is not oriented by $\lfloor \beta_v \rfloor$. As T is thoroughly ordered, this implies $|s| \geq |s_v|$, and hence $s = s_v$ since $|\ |$ is injective.

(ii) The separation s_v distinguishes τ from τ' by definition of v. Since T is thoroughly ordered, any $s \in S$ with $|s| < |s_v|$ has an orientation $\vec{s}$ in $\lfloor \beta_v \rfloor$, which is a subset of both $\tau = \lfloor \beta_\ell \rfloor$ and $\tau' = \lfloor \beta_{\ell'} \rfloor$. Hence $\tau(s) = \tau'(s) = \vec{s}$, so s does not distinguish τ from τ'. □

Lemma 6.5. *N consists of all the separations in S that distinguish two tangles of S optimally.*

Proof. Immediate from Lemma 6.4 and the definition of N. □

Let us call a node v of T *critical* if s_v has an orientation $\vec{s_v}$ that is either co-trivial in $\vec{S}$ or such that $\lfloor \beta_v \rfloor \cup \{\vec{s_v}\}$ contains a triple from $\mathcal{R}(\vec{U})$.

Lemma 6.6. *Tangle nodes of T are never critical.*

Proof. Suppose T has a critical tangle node v. Let $\vec{s_v}$ be as in the definition of its criticality. Let w be the successor of v with $\beta(vw) = \vec{s_v}$, and let $\ell \geq w$ be a tangle leaf of T.

Then $\vec{s_v}$ lies in the $\mathcal{F}$-tangle $\lfloor \beta_\ell \rfloor$. As $\mathcal{F}$ is standard, $\vec{s_v}$ cannot be co-trivial. Hence $\lfloor \beta_v \rfloor \cup \{\vec{s_v}\}$ contains a triple from $\mathcal{R}(\vec{U}) \subseteq \mathcal{F}$. As $\lfloor \beta_v \rfloor \cup \{\vec{s_v}\} \subseteq \lfloor \beta_\ell \rfloor$, this means that ℓ is not a tangle leaf, contrary to its definition. □

Lemma 6.7. *Let u, v be non-leaf nodes of T. If s_u and s_v cross, then either u or v is critical.*

Proof. By Lemma 6.3, we can find $r \in \{s_u, s_v\}$ such that two corners t, t' of s_u and s_v lie on the same side of r and have order $< |r|$. Let $\overleftarrow{r}$ orient r away from these corners. Then these have orientations $\vec{t}, \vec{t'} \leq \vec{r}$. Note that $t, t' \in S$, since $|t|, |t'| < |r|$ and $r \in S = U_k$, as well as $\{\vec{r}, \overleftarrow{t}, \overleftarrow{t'}\} \in \mathcal{R}(\vec{U})$.

[14] See Theorem 3.1 for the definition of $\ell(\tau)$ etc.

Let $r = s_v$, say, and let w be the child of v for which $\beta(vw) = \overleftarrow{r}$. If $\overleftarrow{r}$ is co-trivial in $\vec{S}$, then v is critical as desired, so we assume that $\overleftarrow{r}$ is not co-trivial. Since T is an $\mathcal{F}$-tangle structure tree and v is not a leaf, the set β_v has no subset in $\mathcal{F}$. As $\mathcal{F}$ is standard, this means that β_v, and hence also $\beta_w = \beta_v \cup \{\overleftarrow{r}\}$, has no co-trival elements.

Since t and t' have order $< |r|$ but $s_v = r$, our assumption that T is thoroughly ordered implies that t and t' were ineligible for selection as s_v: they must have orientations in $\lfloor \beta_v \rfloor \subseteq \lfloor \beta_w \rfloor$. Now $\lfloor \beta_w \rfloor$ is consistent by Lemma 2.1, because β_w is consistent by definition of T and has no co-trivial elements. But the only orientations of t and t' that are consistent with $\overleftarrow{r} \in \beta_w$ are $\overleftarrow{t}$ and $\overleftarrow{t'}$, since $\vec{t}, \vec{t'} \leq \vec{r}$. Thus, $\overleftarrow{t}, \overleftarrow{t'} \in \lfloor \beta_v \rfloor$.

This makes $\{\vec{r}, \overleftarrow{t}, \overleftarrow{t'}\} \in \mathcal{R}(\vec{U})$ a witness of the fact that v is critical. □

Lemma 6.8. *N is a tree of the tangles of S. In particular, N is nested.*

Proof. N is nested by Lemmas 6.6 and 6.7. By Lemma 6.5, this makes it a tree of the tangles of S. □

This completes our proof of Theorem 6.2. □

The proof of Theorem 6.2 starts from the premise that we have a thoroughly ordered tangle structure tree for our given separation system. Theorem 3.3 provides this when $\mathcal{F}$ is standard and rich. Both these properties are required in the premise of the theorem; but the richness of $\mathcal{F}$ is never used in the proof. However it follows from the existence of the structure tree that the proof does need, combined with the injectivity of the order function on S [3, Theorem 4.8]. So we may as well require it.

The tangle structure tree in which we find our tree of tangles does not, however, have to be thoroughly ordered: we could work in any tree obtained from the unique thoroughly ordered structure tree provided by Theorem 3.3 by reducing it as in the proof of Theorem 3.4 given in [3]. This is because tangle nodes are never deleted in the reduction process, remain tangle nodes, and their associated separations s_v remain optimal distinguishers of any two tangles such that v is the infimum in T of their tangle leaves. See [3, Section 5] for more analysis.

Theorem 6.2 is our basic tree-of-tangles theorem for robust $\mathcal{F}$-tangles. It applies only to the tangles of a fixed separation system $\vec{S} = \vec{U_k}$ inside some universe $\vec{U}$. The classical tree-of-tangles theorem for profiles [9], however, is more comprehensive (in its more restricted context): it finds a canonical tree of tangles for all the (maximal) tangles in $\vec{S}$, regardless of their order.[15]

We can extend Theorem 6.2 to yield such a more general tree of tangles in $\vec{S}$, too. Our definition of tangle structure trees in this context remains the same as before, except that the notion of tangle leaves is extended to capture all the tangles in $\vec{S}$:

Definition 6.9. An *$\mathcal{F}$-tangle structure tree in $\vec{S}$* is a consistent separation tree T on $\vec{S}$ in which every leaf is either a tangle leaf or forbidden, and for every non-leaf v the set β_v has no subset in $\mathcal{F}$.

Here, a leaf ℓ of T is a *tangle leaf* if $\lfloor \beta_\ell \rfloor$ is a maximal $\mathcal{F}$-tangle in $\vec{S}$. It is *forbidden* if β_ℓ has a subset in $\mathcal{F}$.

Definitions that involve the term 'tangle leaf', such as that of *necessary* edges or nodes of T, adapt accordingly: they now refer to 'tangle leaves' as defined above.

[15] Tangles of S_k are often called 'tangles in $\vec{S}$ of order k'.

Here, then, is our more comprehensive tree-of-tangles theorem. Its tree of tangles displays more tangles than Theorem 6.2 does, namely, all the maximal tangles in $\vec{S}$ rather than just the tangles of S (which are also maximal tangles in $\vec{S}$). However, its premise is also stronger. So neither theorem implies the other.

Theorem 6.10. *Let $\vec{S}$ be a universe of separations with an injective and structurally submodular order function. Let $\mathcal{F}$ be any set that includes $\mathcal{R}(\vec{S})$ and is rich and standard for $\vec{S}_k$ for every $k \in \mathbb{R}$. Then there exists a tree of the tangles in $\vec{S}$.*

Proof. We start by observing that the thoroughly ordered $\mathcal{F}$-tangle structure trees (T_k, r_0, β^k) provided by Theorem 6.2 for the separation systems $\vec{S}_k \subseteq \vec{S}$ are unique by Theorem 3.3. And they are naturally nested: growing from the same root r_0, they satisfy $T_i \subseteq T_j$ and $\beta^i \subseteq \beta^j$ for all $i < j$.

Let n be large enough that $S = S_n$, and let $T := T_n$. Then every non-leaf node v of T is a leaf of some T_k: just choose k big enough that $v \in T_k$, but small enough that $|s_v| \geq k$. As v is a non-leaf node of T, the fact that (T_n, r_0, β^n) is an $\mathcal{F}$-tangle structure tree implies that $\beta^n_v \supseteq \beta^k_v$ has no subset in $\mathcal{F}$. Thus, v is a tangle leaf of T_k in the usual sense of Section 3.

Let T' be obtained from T by deleting any pairs ℓ, ℓ' of forbidden leaves of T that are children of the same node, and let $\beta := \beta^n \restriction E(T')$. Then every leaf of T' is either a forbidden leaf of T, or a tangle leaf ℓ of T_k for some k. In the latter case, its associated tangle τ of S_k is a maximal tangle in $\vec{S}$: it cannot extend to a tangle $\tau' \supsetneq \tau$ of any $S_{k'}$. Indeed, this τ' would be associated with a tangle leaf $\ell' >_r \ell$ of $T_{k'}$, so τ' would include either $\beta_\ell \cup \{\vec{s}_\ell\}$ or $\beta_\ell \cup \{\overleftarrow{s}_\ell\}$. This cannot happen for any $\mathcal{F}$-tangle τ', as both these sets have a subset in $\mathcal{F}$ by the choice of ℓ.

Thus, every leaf ℓ of T' is either forbidden (and also a leaf of T), or it is a tangle leaf in our new sense that $\lfloor \beta_\ell \rfloor$ is a maximal $\mathcal{F}$-tangle in $\vec{S}$. Conversely, for every maximal tangle τ in $\vec{S}$ there is a leaf ℓ of T' such that $\lfloor \beta_\ell \rfloor = \tau$. But, unlike in T, for every non-leaf node v of T' there exists a tangle leaf $\ell > v$ of T'.

Let V be the set of all non-leaf nodes of T' that are not the parent of a forbidden leaf of T. This V is precisely the set of all tangle nodes of T'. Let $N = \{ s_v \mid v \in V \}$. We shall prove that N is our desired tree of tangles: that it is nested, and that it consists of precisely the separations in S that distinguish two maximal $\mathcal{F}$-tangles in $\vec{S}$ optimally.

Let us show first that every two maximal tangles in $\vec{S}$, say a tangle τ_i of S_i and a tangle τ_j of S_j with $i \leq j$, are distinguished optimally by some separation in N. Let us apply Theorem 6.2 with $k := i$ to the tangles τ_i and $\tau_j \cap \vec{S}_i \neq \tau_i$ of S_i. By the proof of the theorem, these tangles are distinguished optimally by s_v for v the infimum of their tangle leaves in T_i. This v lies in V, and s_v also distinguishes τ_i and τ_j optimally as tangles in $\vec{S}$, i.e., as partial orientations of S.

Conversely, let us show that every $s_v \in N$ distinguishes some such pair τ_i, τ_j of tangles in $\vec{S}$ optimally. As v lies in V, it has two children in T', neither of which is a forbidden leaf of T. Since all forbidden leaves of T' are also forbidden leaves of T, the definition of T' implies that v is the infimum of some tangle leaves ℓ_i and ℓ_j of T' associated with tangles τ_i of S_i and τ_j of S_j. By our earlier arguments, these are maximal tangles in $\vec{S}$, and s_v distinguishes them optimally.

To see that N is nested, consider distinct nodes $u, v \in V$ in T. Suppose s_u and s_v cross. By Lemma 6.3, we can find $r \in \{s_u, s_v\}$ such that two corners t, t' of s_u and s_v lie on the same side of r and have order $< |r|$. Let us assume that $r = s_v$. As earlier in our proof, v is a tangle node in some T_i. But now v is critical

in T_i, by the proof of Lemma 6.7. (Note that this proof no longer considers u after identifying v by Lemma 6.3, exactly as we did here. This is important, since in our context u may not be a node of T_i.) This contradicts Lemma 6.6 applied to T_i.[16] □

Let us summarize the details from our proof of Theorem 6.10:

Corollary 6.11. *Let $\vec{S}$ be a universe of separations with an injective and structurally submodular order function. Let $\mathcal{F}$ be any set that includes $\mathcal{R}(\vec{S})$ and is rich and standard for $\vec{S}_k$ for every $k \in \mathbb{R}$. Let T' be obtained from the unique thoroughly ordered $\mathcal{F}$-tangle structure tree (T, r, β) for $\vec{S}$ by deleting any pairs ℓ, ℓ' of forbidden leaves of T that are children of the same node. Let V be the set of tangle nodes of T'.*

Then $N = \{ s_v \mid v \in V \}$ is a tree of the tangles in $\vec{S}$. It is nested and consists of all the separations in S that distinguish two maximal $\mathcal{F}$-tangles in $\vec{S}$ optimally.

7. Discovering riches by zooming in: how to make consistent subsets more efficient

When we compared our new tangle-tree duality results from Section 4 with the classical such theorems, such as Theorem 4.3 [11], we observed two things. One was that our dichotomy Theorem 4.1 is more general than that from [11] since, unlike there, our sets $\mathcal{F}$ of forbidden subsets need not be stars of separations. Due to this greater generality, however, our witnesses for the non-existence of tangles, though 'tree-like' in the sense that they are $\mathcal{F}$-trees, do not impose a tree structure on whatever our separation system separates (e.g., a graph or dataset), as Theorem 4.3 from [11] does.

For the case that $\mathcal{F}$ does consist of stars we went on to prove that our results imply a tangle-tree duality theorem just like Theorem 4.3, albeit from a different premise. At the end of Section 4 we remarked that, since our premise (that $\mathcal{F}$ is rich) differs so fundamentally from the premise in [11] (that $\vec{S}$ is $\mathcal{F}$-separable), it is unlikely that our results could imply those of [11] formally.

Our aim in this section is to show that our premise (and hence, tangle-tree duality) does, however, follow from a slight strengthening of $\mathcal{F}$-separability, one that is actually verified and used when the duality theorem of [11] is applied [10]. As a result, we shall be able to deduce all the applications of [11] to concrete structures of interest [10] from our results. This will be done in [2].

Assuming a structurally submodular injective order function on $\vec{S}$, we shall prove that $\mathcal{F}$ is rich as soon as it is *closed under shifting*. We can then apply our Corollary 4.6 to obtain the same tangle-tree dichotomy as offered by Theorem 4.3, which also makes this assumption about $\mathcal{F}$. In fact, our notion of 'closed under shifting' will be weaker than that in [11], making our result stronger. While [11] makes no formal submodularity assumption, it assumes a consequence of it (separability of $\vec{S}$), which in practice is always established by assuming submodularity [10, Lemma 3.4].

Our strategy for proving, from assumptions similar to those made in [11], that some given $\mathcal{F}$ is rich, is as follows. Recall that $\mathcal{F}$ is *rich* if every consistent orientation τ of S that has a subset in $\mathcal{F}$ also has an efficient such subset: one that is 'maximally focused' in the sense that none of its elements is eclipsed by another separation in τ. Our strategy will be to start from an arbitrary set $\sigma \in \mathcal{F}$, and

[16] We cannot apply the two lemmas to T' directly because T', unlike T and T_i, need not be a tangle structure tree of any $\vec{S'}$ with $S' \subseteq S$.

then to make it more focused by iteratively 'zooming in' until it is efficient.

For intuition, think of any consistent orientation τ of $\vec{S}$ as pointing towards some place in a dataset V that S may be separating. If τ is a tangle, this could be a cluster; if not, it could be the (possibly empty) interior $\bigcap\sigma$ of a star $\sigma \in \mathcal{F}$. For two oriented separations $\vec{r}, \vec{s}$ in τ, think of $\vec{s} > \vec{r}$ as expressing that $\vec{s}$ points towards r and, beyond it, to whatever place $\vec{r}$ points to in V. Thus, while $\vec{s}$ and $\vec{r}$ point to the same place, $\vec{r}$ is closer to it. Hence if $\vec{r}$ eclipses $\vec{s}$, then replacing $\vec{s}$ with $\vec{r}$ in some subset σ of τ – e.g., one in $\mathcal{F}$, or a set of elements whose closure is a tangle of S – makes σ more focused on the place to which τ points.

Our strategy for making a given set $\sigma \in \mathcal{F}$ more efficient will thus be to replace, iteratively, an element of σ with another element of τ that eclipses it. The rest of σ will either be left unchanged (at the start of this section) or 'shifted' to one side of the eclipsing separation, as in [11] (later in this section).

Let us call a collection $\mathcal{F}$ of subsets of $\vec{S}$ *closed under eclipsing* in $\tau \subseteq \vec{S}$ if any set σ' obtained from a subset $\sigma \in \mathcal{F}$ of τ by replacing some $\vec{s} \in \sigma$ with a separation $\vec{r} \in \tau$ that weakly eclipses $\vec{s}$ will also lie in $\mathcal{F}$. Note that σ' will be strictly smaller than σ in the natural extension of our partial ordering $\leq$ on $\vec{S}$ to $2^{\vec{S}}$, in which $\sigma \geq \sigma'$ means that there exists a map $f\colon \sigma \to \sigma'$ such that $\vec{s} \geq f(\vec{s})$ for all $\vec{s} \in \sigma$.

Lemma 7.1. *If $\vec{S}$ is an ordered separation system and $\mathcal{F} \subseteq 2^{\vec{S}}$ is closed under eclipsing in every consistent orientation of S, then $\mathcal{F}$ is rich for $\vec{S}$.*

Proof. To prove that $\mathcal{F}$ is rich, let τ be any consistent orientation of S that has a subset in $\mathcal{F}$. We claim that every set σ that is minimal in $\mathcal{F} \cap 2^\tau \neq \emptyset$ with respect to our partial ordering on $2^{\vec{S}}$ is strongly efficient in τ.

Indeed, suppose σ is not strongly efficient in τ. Then there exists an $\vec{r} \in \tau$ that weakly eclipses some $\vec{s} \in \sigma$. Let $\sigma' := (\sigma \setminus \{\vec{s}\}) \cup \{\vec{r}\}$. Then $\sigma > \sigma' \in \mathcal{F}$, since $\mathcal{F}$ is closed under eclipsing in τ. This contradicts the minimality of σ, since $\sigma' \subseteq \tau$. □

We now introduce our second way of making a set $\sigma \in \mathcal{F}$ more focused, which is essentially the shifting operation from [10, 11]. Let $\vec{s} \in \vec{S}$ be non-trivial in $\vec{S}$ and non-degenerate. Write $S_{\leq \vec{s}}$ for the set of all $t \in S$ to which $\vec{s}$ points, those with an orientation $\vec{t} \leq \vec{s}$. Given any $\vec{r} \leq \vec{s}$ in $\vec{S}$, define the $\vec{S}_{\leq \vec{s}} \to \vec{U}$ map[17]

$$f{\downarrow}^{\vec{s}}_{\vec{r}}\,(\vec{t}) := \begin{cases} \vec{t} \wedge \vec{r} & \text{if } \overleftarrow{s} \neq \vec{t} \leq \vec{s}; \\ (\overleftarrow{t} \wedge \vec{r})^* & \text{if } \overleftarrow{s} \neq \overleftarrow{t} \leq \vec{s}. \end{cases}$$

It is easy to see that such *shifting maps* preserve the partial ordering between their arguments; in particular, they map stars to stars [11]. We say that a separation $\vec{r} \leq \vec{s}$ in $\vec{S}$ *emulates* $\vec{s}$ in $\vec{S}$ if our shifting map has its image inside $\vec{S}$, i.e., if $\vec{t} \wedge \vec{r} \in \vec{S}$ for every $\vec{t} \in \vec{S}$ with $\overleftarrow{s} \neq \vec{t} \leq \vec{s}$.

Let us show that shifting stars can reduce their *order*, the sum of the orders of their elements under a given order function on our separation system.

Lemma 7.2. *Let $\vec{U}$ be a universe of separations with a structurally submodular order function, and let $S = U_\ell$ for some $\ell \in \mathbb{R}$. Let τ be a consistent orientation of S. Assume that some element $\vec{s}$ of a star $\sigma \subseteq \tau$ is non-trivial in $\vec{S}$ and eclipsed by some $\vec{r} \in \tau$. Then $\vec{r}$ can be chosen so that it emulates $\vec{s}$ in $\vec{S}$ and $\sigma' := f{\downarrow}^{\vec{s}}_{\vec{r}}\,(\sigma)$ is a star in τ of lower order than σ.*

[17]Note that, as $\vec{s}$ is non-trivial, $\vec{t} \leq \vec{s}$ and $\overleftarrow{t} \leq \vec{s}$ cannot both hold unless $s = t$, in which case $\vec{s} \wedge \vec{r} = \vec{r}$ (rather than its inverse) is chosen as the image of $\vec{s}$, and $\overleftarrow{r}$ as the image of $\overleftarrow{s}$.

Proof. Choose $\vec{r}$ with $|r|$ minimum, and subject to this maximal under the partial ordering on $\vec{S}$. Let us show first that $\vec{r}$ emulates $\vec{s}$ in $\vec{S}$.

We have to show that $\vec{r} \wedge \vec{t} \in \vec{S}$ for any $\vec{t} \in \vec{S}$ with $\overleftarrow{s} \neq \vec{t} \leq \vec{s}$, i.e., that $|\vec{r} \wedge \vec{t}| < \ell$. By the structural submodularity of our order function we have $|\vec{r} \wedge \vec{t}| \leq |t| < \ell$ as desired if $|\vec{r} \vee \vec{t}| > |r|$, so let us assume that $|\vec{r} \vee \vec{t}| \leq |r|$. Then $\vec{r} \vee \vec{t} \in \vec{S}$; let us show that, in fact, $\vec{r} \vee \vec{t} \in \tau$.

As $\vec{r} \vee \vec{t} \geq \vec{r} \in \tau$, the consistency of τ implies $\vec{r} \vee \vec{t} \in \tau$ as desired, unless the unoriented separation underlying $\vec{r} \vee \vec{t}$ is r. If $\vec{r} \vee \vec{t} = \vec{r}$, we are done since $\vec{r} \in \tau$, so let us assume that $\overleftarrow{r} = \vec{r} \vee \vec{t} \geq \vec{r}$. Then $\vec{s} \geq \vec{r} \vee \vec{t} = \overleftarrow{r}$, since $\vec{s} \geq \vec{r}$ as well as $\vec{s} \geq \vec{t}$ by assumption, which makes $\vec{s}$ trivial witnessed by r, contrary to assumption, unless $s = r$. But if $\vec{s} = \vec{r}$ we have $\vec{r} \vee \vec{t} = \vec{s} \vee \vec{t} = \vec{s} \in \tau$, while if $\vec{s} = \overleftarrow{r}$ we have $\vec{r} \vee \vec{t} = \overleftarrow{r} = \vec{s} \in \tau$. This completes our proof that $\vec{r} \vee \vec{t} \in \tau$.

As noted, we have $\vec{s} \geq \vec{r} \vee \vec{t}$, as well as $|\vec{r} \vee \vec{t}| \leq |r| < |s|$, since $\vec{r}$ eclipses $\vec{s}$. So $\vec{r} \vee \vec{t}$ eclipses $\vec{s}$, and was thus a candidate for $\vec{r}$. As $|\vec{r} \vee \vec{t}| \leq |r|$, this contradicts the choice of $\vec{r}$ unless $\vec{r} \vee \vec{t} = \vec{r}$. But in that case $\vec{r} \geq \vec{t}$, giving $\vec{r} \wedge \vec{t} = \vec{t} \in \vec{S}$ as desired. This completes our proof that $\vec{r}$ emulates $\vec{s}$.

As shifting preserves the partial ordering on $\vec{S}$, we know that σ' is a star. Let us show that it has lower order than σ. Since $f\downarrow_{\vec{r}}^{\vec{s}}(\vec{s}) = \vec{r}$ and $|r| < |s|$, it suffices to show that $\left|f\downarrow_{\vec{r}}^{\vec{s}}(\vec{t})\right| = |\vec{r} \wedge \vec{t}| \leq |t|$ for any other elements $\overleftarrow{t}$ of σ. But this follows from structural submodularity and the choice of $\vec{r}$, as earlier, unless $\vec{r} \wedge \vec{t} = \vec{t}$. But in that case we also have $|\vec{r} \wedge \vec{t}| = |t|$ as desired.

It remains to show that $\sigma' \subseteq \tau$. We have $f\downarrow_{\vec{r}}^{\vec{s}}(\vec{s}) = \vec{r} \in \tau$ by explicit assumption about $\vec{r}$. For any $\overleftarrow{t} \in \sigma$ other than $\vec{s}$, we have $\vec{s} \geq \vec{t}$, since σ is a star, and hence $f\downarrow_{\vec{r}}^{\vec{s}}(\overleftarrow{t}) = (\vec{t} \wedge \vec{r})^* = \overleftarrow{t} \vee \overleftarrow{r} \geq \overleftarrow{t} \in \tau$ by the definition of $f\downarrow_{\vec{r}}^{\vec{s}}$. This implies $f\downarrow_{\vec{r}}^{\vec{s}}(\overleftarrow{t}) \in \tau$ by the consistency of τ unless $f\downarrow_{\vec{r}}^{\vec{s}}(\overleftarrow{t}) = \vec{t}$, so let us show that this cannot be the case.

If it is, then we have shown that $\vec{t} = f\downarrow_{\vec{r}}^{\vec{s}}(\overleftarrow{t}) \geq \overleftarrow{t}$. Then $\vec{s} \geq \vec{t} \geq \overleftarrow{t}$, which contradicts our assumption that $\vec{s}$ is non-trivial unless $s = t$. But in that case we have $\vec{s} = \vec{t}$, since $\vec{s} \neq \overleftarrow{t}$ as these are distinct elements of σ. But if $\vec{s} = \vec{t}$ then $\overleftarrow{r} = f\downarrow_{\vec{r}}^{\vec{s}}(\overleftarrow{s}) = f\downarrow_{\vec{r}}^{\vec{s}}(\overleftarrow{t}) = \vec{t} = \vec{s}$, which contradicts the fact that $|r| < |s|$. □

Call a set $\mathcal{F} \subseteq 2^{\vec{S}}$ of stars *closed under shifting*[18] in $\tau \subseteq \vec{S}$ if for every $\sigma \in \mathcal{F} \cap 2^{\tau}$ and every $\vec{r} \in \tau$ that weakly eclipses and emulates some non-trivial $\vec{s} \in \sigma$ in $\vec{S}$ we have $f\downarrow_{\vec{r}}^{\vec{s}}(\sigma) \in \mathcal{F}$.

Lemma 7.3. *Let $\vec{U}$ be a universe of separations with a structurally submodular injective order function, and let $S = U_\ell$ for some $\ell \in \mathbb{R}$. Let $\mathcal{F}$ be a set of stars of non-trivial separations in $\vec{S}$ that is closed under shifting in $\vec{S}$.*[18] *Then $\mathcal{F}$ is rich for $\vec{S}$.*

Proof. To prove that $\mathcal{F}$ is rich, consider any consistent orientation τ of S that includes some $\sigma \in \mathcal{F}$. Choose such a star σ in τ of minimum order. We shall show that σ is strongly efficient in τ.

If not, there exists an $\vec{r} \in \tau$ that weakly eclipses some $\vec{s} \in \sigma$. Since our order function is injective, any such $\vec{r}$ in fact eclipses $\vec{s}$. By Lemma 7.2, we can choose $\vec{r}$ so that it emulates $\vec{s}$ in $\vec{S}$ and $\sigma' := f\downarrow_{\vec{r}}^{\vec{s}}(\sigma)$ is a star in τ of lower order than σ. As $\mathcal{F}$ is closed under shifting we have $\sigma' \in \mathcal{F}$, which contradicts our choice of σ. □

[18] Our definition of 'closed under shifting' is weaker than that in [10,11], making our subsequent results stronger. We can even strengthen them further: as is easily checked, we can weaken in both Lemma 7.3 and Theorem 7.4 the assumption that $\mathcal{F}$ is closed under shifting in $\vec{S}$ to assuming that $\mathcal{F}$ is closed under shifting in every consistent orientation of S. We omitted this for simplicity.

Theorem 7.4. *Let $\vec{U}$ be a universe of separations with either a submodular or an injective structurally submodular order function. Let $S = U_\ell$ for some $\ell \in \mathbb{R}$. Assume that $\vec{S}$ has no trivial elements. Let $\mathcal{F}$ be a set of stars of separations in $\vec{S}$ that is closed under shifting in $\vec{S}$.*[18] *Then exactly one of the following assertions holds:*

(i) *there exists an $\mathcal{F}$-tangle of S;*

(ii) *there exists an S-tree over $\mathcal{F}$.*

Proof. Recall from the proof of Lemma 5.1 that if S has the form of $S = U_\ell$ with respect to a given order function o, it also has this form with respect to any refinement of o. By Lemma 5.6 we may therefore assume that our order function on U is injective and structurally submodular. By Lemma 7.3, $\mathcal{F}$ is rich for $\vec{S}$. The theorem now follows from Corollary 4.6. □

While Theorem 7.4 is probably to be the most widely applicable tangle-tree duality theorem to date, it may appear restrictive that S must have the form of U_ℓ, and that we need a submodular order function on U. The original tangle-tree duality theorem from [11], Theorem 4.3, makes no such assumptions. However, it makes another assumption, one not made in Theorem 7.4: that $\vec{S}$ is $\mathcal{F}$-separable.

This is not a coincidence: $\mathcal{F}$-separability is essentially a technical summary of exactly those consequences of our two more natural assumptions, submodularity and $S = U_\ell$, that are technically needed in the proof of Theorem 4.3.

Finally, the premise of Theorem 7.4 differs from that of Theorem 4.3 in that it requires $\vec{S}$ not to have trivial elements. We need this for our conversion of $\mathcal{F}$-trees into S-trees over $\mathcal{F}$ in Theorem 4.4. But it entails no loss of generality: the trivial separations of any separation system lie in all consistent orientations, so we can just delete them from our given system and add them back later if desired [4, 5].

References

[1] Sandra Albrechtsen, *Refining trees of tangles in abstract separation systems: Inessential parts (extended version)*, Combinatorial Theory **4(1)** (2024; arXiv:2302.01808v1).

[2] Hanno von Bergen, *Tangle structure tree applications*, in preparation.

[3] H.v. Bergen and R. Diestel, *Tangle structure trees*, arXiv:2601.15870.

[4] R. Diestel, *Abstract separation systems*, Order **35** (2018), 157–170; arXiv:1406.3797.

[5] ———, *Tree sets*, Order **35** (2018), 171–192.

[6] ———, *Tangles: A structural approach to artificial intelligence in the empirical sciences*, Cambridge University Press, 2024. e-book: `tangles-book.com/book/`.

[7] ———, *Graph theory* (6th edition), Springer-Verlag, 2025.
Electronic edition available at `http://diestel-graph-theory.com/`.

[8] R. Diestel, J. Erde, and D. Weißauer, *Structural submodularity and tangles in abstract separation systems*, J. Comb. Theory Ser. A **167C** (2019), 155–180; arXiv:1805.01439.

[9] R. Diestel, F. Hundertmark, and S. Lemanczyk, *Profiles of separations: in graphs, matroids, and beyond*, Combinatorica **39** (2019), no. 1, 37–75; arXiv:1110.6207.

[10] R. Diestel and S. Oum, *Tangle-tree duality in graphs, matroids and beyond*, Combinatorica **39** (2019), 879–910. arXiv:1701.02651.

[11] ———, *Tangle-tree duality in abstract separation systems*, Advances in Mathematics **377** (2021), 107470; arXiv:1701.02509.

[12] N. Robertson and P.D. Seymour, *Graph minors. X. Obstructions to tree-decomposition*, J. Comb. Theory, Series B **52** (1991), 153–190.